Journal of Artificial Intelligence and Technology, (Ahead of Print)
https://doi.org/10.37965/jait.2026.1067

RESEARCH ARTICLE

ConReNet: A Modular Context-Aware Remedy Recommendation Network for Plant Disease Management

Jayashree R and Kusuma Kumari B M
Department of Studies and Research in Computer Application, Jnanasiri Campus, Tumkur University, Bidrakatte-572118, Karnataka, India

(Received 08 December 2025; Revised 13 March 2026; Accepted 19 May 2026; Published online 11 June 2026)

Abstract: Plant diseases remain a major challenge for global agriculture, causing significant yield losses and economic damage each year. Although recent advances in artificial intelligence have improved disease detection accuracy, most existing systems are limited to diagnostic tasks and rely solely on image data, overlooking critical contextual factors such as crop stage, disease severity, and environmental conditions. To address these limitations, this paper proposes Context-Aware Remedy Recommendation Network (ConReNet), a lightweight, modular, end-to-end framework designed to generate crop- and context-specific remedy recommendations. ConReNet consists of three core components: (i) the Contextual Feature Interaction Encoder (CFIE), which captures complex dependencies among crop, disease, and environmental variables; (ii) the Sparse Contextual Fusion Module (SCFM), which efficiently aggregates heterogeneous contextual information; and (iii) the Context-Aware Remedy Ranking Network (CARRN), which produces ranked remedy recommendations. Synthetic datasets simulating realistic crop–disease–remedy interactions were constructed to evaluate the framework. Experimental results demonstrate that ConReNet effectively learns contextual relationships and delivers accurate recommendations, establishing its potential as a scalable solution for intelligent disease management. Its modular design enables future integration with image-based models and IoT data streams, making it suitable for real-world deployment in precision agriculture.

Keywords: ConReNet framework; Context-Aware Remedy Recommendation; Contextual Feature Interaction Encoder (CFIE); precision agriculture; Sparse Contextual Fusion Module (SCFM)

I. INTRODUCTION

Plant diseases represent one of the most significant threats to global agriculture, affecting crop yield, food quality, and economic stability across both developed and developing regions. These diseases, caused by a variety of pathogens including fungi, bacteria, viruses, and nematodes, can spread rapidly under favorable environmental conditions, leading to substantial losses if not identified and treated promptly. Globally, plant diseases are responsible for up to 40% of annual crop losses, posing a serious threat to food security and farmer livelihoods. For farmers, especially those in resource-limited settings, early detection and appropriate management of plant diseases remain a persistent challenge due to limited access to expert knowledge, diagnostic tools, and real-time advisory services [1–3].

Traditional methods for disease identification often rely on visual inspection, expert consultation, and manual referencing of agronomic literature, which may be time-consuming, inconsistent, and inaccessible at scale. Moreover, even when accurate classification is achieved, determining the correct remedy—such as the selection of appropriate pesticides, cultural practices, or biological treatments—requires additional domain expertise and contextual awareness, such as crop type, disease severity, and environmental constraints. In this context, the integration of computational tools into plant disease management has opened new possibilities for precision agriculture. By combining automated disease classification with intelligent recommendation systems, it becomes feasible to deliver timely, localized, and actionable suggestions to farmers and agricultural advisors. Such systems can enhance decision-making by linking observed disease symptoms with curated remedy options, thereby improving crop health outcomes while minimizing the misuse of agrochemicals [4–7]. In modern agriculture, the early detection and accurate treatment of plant diseases are critical to ensuring crop health, yield stability, and sustainable farming practices. While advancements in machine learning have enabled high-accuracy classification of plant diseases from images and field data, the translation of these classifications into actionable remedy recommendations remains an unresolved challenge. Existing systems often stop at disease detection and fail to provide context-aware, trustworthy, and explainable treatment options, especially when environmental, crop-specific, and regional agronomic variables are involved. Moreover, most farmers—particularly in developing regions—lack access to expert consultation or timely guidance, leading to over-reliance on general-purpose chemical treatments. This not only increases production costs but also contributes to pesticide resistance, soil degradation, and environmental harm [8,9]. While most existing systems rely predominantly on image-based disease detection, they often ignore crucial contextual information such as crop growth stage, disease severity, and environmental conditions. Integrating such tabular data can significantly enhance recommendation quality and make predictions more interpretable. There is a pressing need for an

Corresponding author: Jayashree R (e-mail: jayashreer@tumkuruniversity.ac.in).

integrated system that not only classifies plant diseases accurately but also recommends targeted remedies by incorporating domain knowledge, field-level data, and explainable AI models. Such a system should support transparent, adaptive, and data-driven decision-making, thereby bridging this critical gap. To address these challenges, this paper proposes an Explainable and Confidence-Aware Remedy Recommendation Framework that bridges the gap between disease detection and actionable treatment. The framework combines contextual feature encoding, sparse feature fusion, remedy ranking, and interpretability mechanisms to deliver accurate, transparent, and context-driven recommendations. Unlike conventional approaches, it provides not only diagnostic outputs but also interpretable, ranked remedy suggestions that align with real-world agronomic practices.

The agricultural sector is increasingly challenged by rising incidences of plant diseases, driven by climate variability, monoculture practices, and evolving pathogen resistance [10–13]. Early disease detection has benefited significantly from advances in artificial intelligence, particularly image-based classification models. However, the journey from disease diagnosis to actionable field treatment remains incomplete. Farmers, especially in resource-constrained or remote regions, often struggle to translate technical disease names into practical actions—such as choosing the right pesticide, applying organic alternatives, or adjusting agronomic practices. This disconnect between disease identification and remedy selection is a critical bottleneck in realizing the full potential of AI in precision agriculture. Existing advisory systems often lack context awareness, are difficult to interpret, or provide generic recommendations that overlook crop-specific, environmental, or regulatory constraints. Furthermore, the absence of explainability in many AI-driven systems erodes user trust, limiting adoption at scale. To address these gaps, this work proposes a modular, end-to-end remedy recommendation framework designed to integrate contextual information and deliver accurate, crop-specific treatment suggestions. The key contributions of this work are summarized as follows:

End-to-End Contextual Remedy Recommendation Framework: We propose a complete pipeline that combines contextual feature encoding, sparse feature fusion, and remedy ranking to bridge the gap between disease detection and actionable recommendations. Contextual Feature Interaction Encoder (CFIE): CFIE captures complex relationships between crop, disease, and environmental variables, providing rich feature representations for downstream processing. Sparse Contextual Fusion Module (SCFM): SCFM integrates heterogeneous contextual features in a computationally efficient manner, enhancing the representation quality for remedy ranking. Context-Aware Remedy Ranking Network (CARRN): A TabNet-based ranking network is employed to generate remedy recommendations based on contextual inputs, enabling accurate and scalable decision support. Lightweight and Modular Design: The proposed architecture is lightweight, making it suitable for future integration with image-based models or IoT systems for real-world deployment.

II. RELATED WORK

The integration of machine learning and artificial intelligence in agricultural diagnostics has led to substantial progress in automated plant disease detection and remedy recommendation systems. This section reviews state-of-the-art contributions that directly align with our objective of developing an intelligent system for disease classification and remedy guidance. Reference [14] presented an overview of how AI and ML algorithms are reshaping traditional agricultural practices by providing automated systems for disease classification, analysis, and prevention. Their work emphasizes increased productivity and reduced crop losses by leveraging data-driven diagnostics. Reference [15] introduced a comprehensive AI system that integrates CNNs and traditional machine learning models to address the limitations of manual disease identification. The study highlights how advanced architectures can provide early-stage detection and support downstream remedy recommendations, making the system scalable across environments. Reference [16] provided a broad review of AI-based plant disease tracking, exploring tools that integrate computer vision, satellite imagery, and IoT. Their emphasis on mobile apps and real-time prediction aligns with modern systems that provide remedy advice based on sensor and image data. Reference [17] conducted a comparative study on machine learning and deep learning models for disease diagnosis. Their analysis showed that deep learning models provide superior performance, especially in complex plant environments, enabling more accurate treatment guidance. Reference [18] reviewed decision support systems based on machine learning and remote sensing, showing how they help farmers choose appropriate treatments and localize disease intervention areas, enabling sustainable and cost-effective agriculture. Reference [19] highlighted the speed and scalability advantages of AI in real-time plant disease detection. Their work demonstrates that AI models can boost agricultural productivity by triggering early remedy actions, reducing the window for disease progression. Reference [20] introduced Agri Shield, a CNN-powered platform capable of diagnosing 20 + plant diseases with 93% accuracy and delivering sustainable, customized remedy recommendations. The paper stands out for combining performance with eco-conscious decision-making. Reference [21] proposed a three-stage AI system involving preprocessing, classification, and supplementary modules for disease detection and fertilizer recommendation. The system targets six plant families and 39 disease classes, demonstrating strong generalization capabilities. Reference [22] explored an end-to-end plant care system integrating diagnosis, remedy recommendation, and local service suggestions. Their solution includes image analysis, chatbot support, and Google Maps integration for accessible decision-making tools. Reference [23] provided a detailed survey of plant disease datasets, ML models, and application challenges. The work identifies limitations like data imbalance and scalability of remedy recommendation systems, guiding future model improvements. Reference [24] conducted an evaluation of various image processing and ML frameworks for disease classification. The comparative results inform model selection for reliable remedy generation, highlighting areas for improvement in interpretability. Reference [25] proposed a mobile application that combines CNN-based segmentation and disease classification with a recommendation engine for treatment and prevention. It showcases the potential of real-time, accessible AI tools for precision agricultural interventions. In [26], PlantCareNet, an automated end-to-end plant disease diagnostic system, was developed that combines deep learning with expert knowledge for precise detection and treatment guidance. It employs a CNN with Dense-100 and Dense-35 layers for accurate disease classification from leaf images. The system then provides both automated and expert-guided recommendations for timely and personalized intervention. In [27–30], AIDoctor Plant a lightweight, few-shot learning framework implemented as an Android application for plant disease diagnosis and remedy recommendation. It supports both English and Hindi interfaces and incorporates disease severity to provide tailored disease management advisories.

III. PROPOSED METHODOLOGY

Agriculture is leaving behind traditional ways of doing things and moving toward decisions based on data and intelligence. AI is being utilized to improve the precision agriculture for disease detection and recommendation of remedy toward it. Figure 1 shows the proposed model and AI-based remedy recommendation for multicrop disease.

Problem Formulation. The remedy recommendation task in Context-Aware Remedy Recommendation Network (ConReNet) is formulated as a top-k ranking problem. Given an input context vector:

$$z = (crop\_type,\ disease\_class,\ severity,\ growth\_stage,\ region,\ confidence)$$

the objective is to produce a ranked list of remedy options ordered by predicted relevance:

$$R = \{r_1, r_2, \ldots, r_k\},\ r_i \in \{chemical,\ organic,\ integrated\}$$

The system employs a confidence-based abstention mechanism: if the maximum predicted probability falls below a threshold τ, the system abstains from making a recommendation, prioritizing reliability over coverage. This formulation is distinct from multi-label classification and traditional learning-to-rank approaches. The modular pipeline processes input as:

$$Z \rightarrow CFIE(Z) \rightarrow SCFM(CFIE(Z)) \rightarrow CARRN(SCFM(CFIE(Z))) \rightarrow R$$

Each module is independently trainable and replaceable, enabling seamless future integration with image-based classifiers or IoT sensor data.

Architecture Specification. The detailed architecture of ConReNet is as follows:

Module	Architecture Details
CFIE	2-layer transformer encoder, 4 attention heads, embedding dim d = 64, ReLU-activated FFN with hidden dim 128
SCFM	Sparse masked attention with learnable Q, K, V projection matrices; binary adaptive sparsity mask for computational efficiency
CARRN	TabNet-based architecture, N_steps = 3, relaxation factor $\gamma = 1.5$, sequential attention for remedy ranking
Pipeline	Tabular context → CFIE → SCFM → CARRN → Ranked remedies with confidence scores

The complete pipeline accepts tabular contextual features as input and outputs ranked remedy recommendations with associated confidence scores.

Figure 1 illustrates a high-level AI-driven agricultural decision support pipeline, consisting of data acquisition, data preparation, data processing, decision-making, and service delivery stages. The proposed framework is situated primarily within the Data Processing and Decision Making phases, highlighted in red. At this stage, contextual tabular data—derived from historical records, farmer experience, and environmental factors—are processed through the CFIE and the SCFM to capture complex feature interactions. These

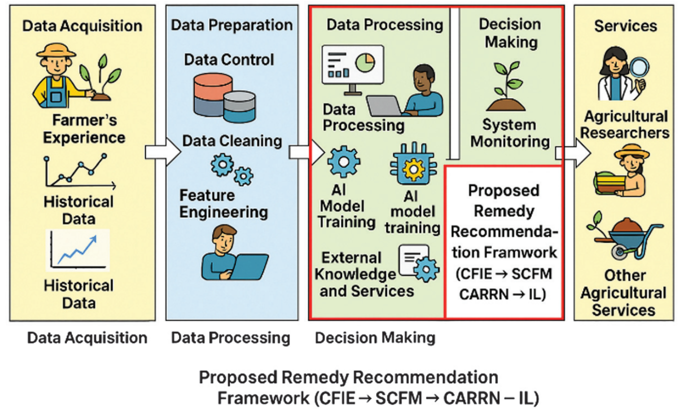

Fig. 1. Proposed model.

representations are then passed to the CARRN, which leverages CARRN to generate ranked remedy recommendations. Finally, the Interpretability Layer (IL) integrates intrinsic feature masks and post hoc explanation techniques (e.g., SHAP, LIME) to provide transparent, human-understandable explanations for each recommendation. This structured integration ensures that remedy suggestions are not only accurate but also auditable and aligned with domain knowledge, facilitating practical decision-making by farmers, agronomists, and other agricultural service stakeholders.

Figure 2 presents the model training and evaluation workflow adopted in this study. The process begins with data preprocessing, including feature selection, handling missing values, and encoding categorical variables to prepare the contextual agricultural dataset. After splitting the data into training and testing sets, the CFIE

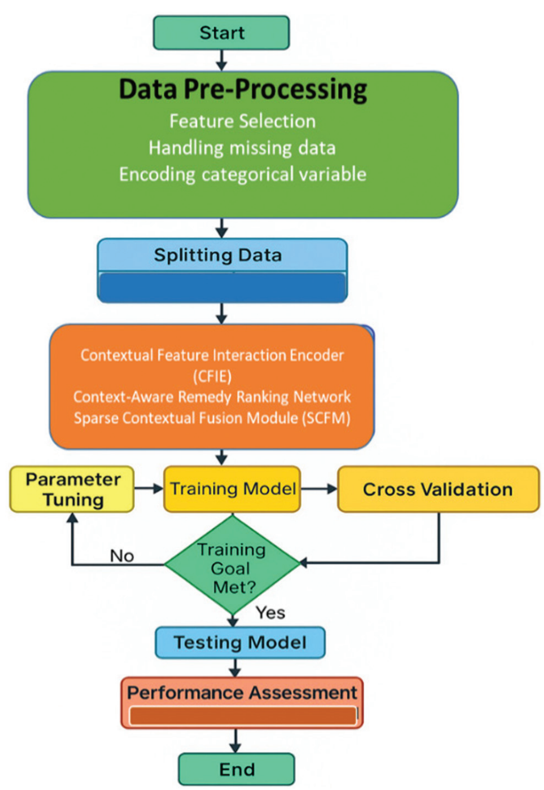

Fig. 2. Proposed model workflow considering the data.

extracts complex cross-feature dependencies, while the SCFM integrates heterogeneous contextual signals. These representations are passed to the CARRN to learn remedy ranking patterns based on disease, crop, and environmental features. Model training is followed by cross-validation to assess generalization and parameter tuning to optimize performance. This loop continues until the training goal is achieved. Once training converges, the model undergoes testing and performance assessment, using relevant evaluation metrics to ensure robustness and accuracy of remedy recommendations. This structured workflow enables the development of a well-validated, interpretable, and domain-aligned recommendation model.

A. DATA PREPROCESSING

Data preprocessing includes Normalization, dealing with missing values, and encoding approaches are all parts of data pretreatment that make sure the data work with DL models. This multi-modal dataset offers a solid basis for recording many connections that exist within the agricultural ecosystem. The steps listed below are part of data preprocessing: The first step is to clean the data, which means finding and fixing any missing or incomplete numbers. In real-world agricultural data, it is usual for entries to be missing because sensors do not work right or records are not always accurate. Mean or median imputation, for example, is used to fill in these gaps, depending on what kind of feature it is. Normalization is also used to bring numerical data like temperature, moisture, and nutrient levels into a regular range, which is usually between 0 and 1. This makes sure that no one characteristic stands out in model learning because of its size. Removing outliers is another important step. Standard deviation thresholds or visual analysis approaches are used to find extreme or uncommon values in the dataset that are produced by bad sensor readings or rare environmental conditions. If possible, these outliers are fixed; if not, they are taken out so they do not affect the model's predictions. The dataset has categorical information as well, including crop labels (rice, wheat, maize, etc.). To turn these categories into numbers, encoding techniques are applied. This is done by employing label encoding or one-hot encoding, which makes the data work with machine learning algorithms that need numbers as input.

B. SPARSE CONTEXTUAL FUSION MODULE (SCFM)

SCFM is a neural network architecture designed to efficiently integrate information from multiple data sources. It was chosen for its effectiveness in managing complex agricultural datasets, which involve overlapping variables, dynamic conditions, and deeply connected factors influencing crop performance and resource utilization. Unlike traditional machine learning models such as Random Forests, Support Vector Machines (SVMs), and Multilayer Perceptron (MLPs), SCFM leverages transformer-based attention mechanisms that can capture long-range dependencies among features. Interdependence of Features and long-range Attention Weather, soil qualities, and records of farmer activity are all examples of different inputs that make up agricultural data. SCFM's attention mechanism does a better job of grasping long-range dependencies than standard models because it properly captures the links between these properties. This leads to predictions that are more accurate and reliable. Fusion of efficient features SCFM has a special feature fusion layer after the self-attention module that cleans up and combines complicated information from several sources. This approach makes it easier for the model to find important patterns, which makes its predictions more accurate and reliable.

C. CONTEXTUAL FEATURE INTERACTION ENCODER (CFIE)

It is processing because they are good at seeing local patterns in fixed spatial positions. This method does not work well with agricultural data, which is usually in Table I form and does not have a regular structure or geographic order. CFIE gets around this problem by adopting a self-attention process that looks at how features like soil moisture, temperature, and humidity are related to each other. First, each feature is changed into a rich, high-dimensional representation. Then, it goes via several transformer encoder layers. These layers help the model learn patterns in the data that are useful and take into account the context. CFIE is better than CNNs or regular fully connected networks for representing the complicated, nonlinear interactions that are often seen in agricultural datasets. It shows that it can generalize well, adapt to new

Algorithm 1. Sparse Weighted Fusion Transformer

Input	$Z \in T^{p*f}$ Y_S, Y_M, Y_X: learnable projection matrices $A \in \{0,1\}^{n*n}$: adaptive sparsity mask Fuse($\cdot$): fusion operator (weighted aggregation) y_{true}: ground-truth labels (training only)
Step 1	**Feature Projections (Input Encoding).** Project inputs to query, key, and value spaces: $S = ZY_S \quad M = ZY_M \quad X = ZY_X$
Step 2	**Sparse Attention Scores.** Compute masked attention logits and normalize: $U = \frac{SM^V}{\sqrt{f_m}} \quad U_{masked} = U \odot C \quad C_{sparse} = softmax(U_{masked})$ The binary mask C suppresses irrelevant connections and enforces sparsity.
Step 3	**Weighted Fusion.** Aggregate values using sparse attention: $H_{fused} = Fuse\ (C_{sparse}, X)$
Step 4	**Transformer Processing (SCFM Layers).** Pass H_{fused} through one or more SCFM blocks (positional encoding, normalization, and feed-forward/attention refinements) to obtain J
Step 5	**Prediction Head.** Map J to task outputs with an MLP head: $A_{pred} = MLP(J)$
Step 6	**Training Objective (if supervised training).** Use cross-entropy for classification or RMSE/MSE for regression; update all learnable parameters via backpropagation.
Return	**Output**

Table I. Contextual Feature Interaction Encoder (CFIE)

Input	Z, Y_G, V, Y_{he}, A
Step 1	Feature embedding: transform the input features into a higher dimension representation through linear transformation $$G(Z) = ZY_G + d_G,$$
Step 2	Transformation of encoder processes the embedded features through the within two transformer encoder layers. each layer consists of multi-head self-attention feed-forward network • Each attention head evaluation $$Att(S,M,X) = softmax\left(\frac{SM^V}{\sqrt{f_m}}\right)$$ • Multi head attention output $$MH(Z) = concat(head_1, \ldots, head_j)Y_q$$ • FFN: after attention the data is transferred through a two-layer FFN
Step 3	Classification Layer: After going through the transformer encoder, the last feature is linked to the output classes: $$R(A\|Z) = softmax(Y_{he}B + d_{he})$$ The predicted class is: $$A = argmaxR(A\|Z)$$
Step 4	Training Process: a. Loss Function (Cross Entropy): For multiclass classification the loss function computed as: $$n = -\sum_{k=1}^{e} a_k logR(a_k\|Z)$$ Optimization algorithm (Adam): Gradient updates are computed using back propagation: $$\theta < -\theta - \vartheta \nabla_\theta n$$
Step 5	Training loop: For each epoch, update the model parameters: $$\theta^{(v+1)} = \theta^v - \vartheta \sum_{k=1}^{P} \nabla_\theta n(Z_k, A_k)$$ **End algorithm**

situations, and handle real-world, noisy sensor data well. Because of this, CFIE is great for things like analyzing soil, planning irrigation, and making decisions about crops.

where Z is the input tabular data of shape (P,F) and is the learning weight matrix for embedding. V is a transformer encoder with self-attention layers. Y_{he} is a fully connected layer for classification, and A is a variable that supervised learning target labels

Feature embedding: Here, Y_G is a learnable weight matrix for the embedding, d_G denotes the bias term, and $G(Z)$ denotes the embedded representation of input in Equation (1). Multi-head self-attention: Equation (2) is learnable weights for query, key, and value. f_m denotes a key dimension in Equation (2). Multi-head attention output: Y_q denotes the output projection matrix in Equation (3). FFN represents the feed-forward network. However, Y_1 and Y_2 represent the weight matrices, d_1 and d_2 denote the biases, and α denotes the ReLU activation function in Equation (4). Classification layer: Here, B denotes the output from the transformer encoder, and Y_{he} and d_{he} denote learnable classification parameter in Equation (5). Loss function C denotes the number of classes and a_k denotes the true label with Adam optimization $\theta < -\theta - \gamma \nabla_\theta n$. Where γ denotes the learning rate; furthermore, training loop is given as in Equation (6). Where Z denotes the output and A denotes the target labels in Equation (7):

$$G(Z) = ZY_G + d_G \tag{1}$$

$$Attn(S,M,X) = softmax\left(\frac{SM^V}{\sqrt{f_m}}\right) \tag{2}$$

$$\begin{gathered} S = G(Z)Y_{SX}M,\ M = G(Z)Y_M,\ X = G(Z)Y_X, Y_M, Y_X \\ MultiHead(Z) = Concat(head_1, \ldots\ldots\ldots, head_j)Y_q \end{gathered} \tag{3}$$

$$FFN(Z) = \alpha(Y_1 Z + d_1)Y_2 + d_2 \tag{4}$$

$$R(A|Z) = softmax(Y_{he}B + d_{he}) \tag{5}$$

$$n = -\sum_{k=1}^{e} a_k logR(Z_k|z) \tag{6}$$

$$\theta^{(v+1)} = \theta^{(v)} - \gamma \sum_{k=1}^{P} \nabla_\theta n(Z_k, A_k) \tag{7}$$

D. CONTEXT-REMEDY RANKING NETWORK + INTERPRETABILITY LAYER AWARE

The following algorithm outlines the prediction and local explanation procedure applied within the CARRN + IL block. Given a new contextual input sample, the trained model predicts the most suitable remedy and then generates a local feature importance explanation using the LIME method. This ensures that each recommendation is accompanied by interpretable evidence, enabling transparent decision-making. The CARRN + IL algorithm integrates predictive intelligence with interpretability by coupling the TabNet-based ranking model with LIME-based local explanation. This allows the system to not only predict optimal remedies based on contextual agricultural data but also explain why a particular remedy was recommended, increasing transparency and trust for agronomists and farmers. Table II shows that Context-Remedy Ranking Network + Interpretability Layer Aware algorithm.

IV. PERFORMANCE EVALUATION

The proposed ConReNet framework effectively bridges the gap between plant disease detection and actionable remedy

Table II. Context-Remedy Ranking Network + Interpretability Layer Aware Algorithm

Input	$z_{new}, Z_{train}, A_{train}, h_{tn}, Lim_{exp}$
Step 1	Preprocessing: Normalize and preprocess z_{new} according on how the training data is spread out.
Step 2	Prediction using Tn Pass z_{new} with the trained Tn model $$R(a\|z_{new}) = h_{tn}(z_{new})$$ Retrieve the predicted class label as $$a = argmax\, R(a\|z_{new})$$
Step 3	Generate LIME exp: Generate different samples Z with z_{new} $$Z = \{z_{new} + \tau_k \| \tau_k \sim P(0,\delta^2)\}$$ Compute model predictions on perturbed sample: $$R(a\|Z) - h_{tn}(Z)$$ Fit a locally weighted linear model (z) to approximate $$h_{tn}{:}i(z)\text{-}\rho_0 + \sum_{k-1}^{f} \rho_k z_k$$
Step 4	Extraction $\varphi_{z_{new}} - (\rho_0, \rho_1, \ldots\ldots\ldots, \rho_0)$
Step 5	Return *pred* and *explain*: output a with feature importance $\varphi_{z_{new}}$ End algorithm

recommendation by leveraging contextual agricultural data. CFIE enhances representation by capturing complex crop–disease–environment interactions, while SCFM efficiently fuses heterogeneous features with minimal overhead. CARRN enables accurate ranking of treatment options, reflecting real-world decision-making scenarios. The modular and lightweight design facilitates integration with image-based models and IoT systems for scalable deployment. However, real-world performance may depend on the availability and quality of contextual data.

A. DATASET DETAILS

In this study, we utilized a synthetically generated multicrop dataset to simulate realistic agricultural conditions for disease detection and remedy recommendation. The dataset encompassed four major crops—Corn (4188 samples), Tomato (3000 samples), Rice (2658 samples), and Sugarcane (1000 samples)—with disease class distributions aligned to observed field prevalence. For Corn, the classes included Healthy (25%), Common Rust (30%), Gray Leaf Spot (25%), and Blight (20%). Tomato comprised five categories: Healthy (20%), Early Blight (25%), Late Blight (25%), Leaf Mold (15%), and Septoria Leaf Spot (15%). Rice included Healthy (25%), Brown Spot (25%), Leaf Blast (30%), and Hispa (20%). Sugarcane data were balanced across Healthy, Red Rot, Mosaic, and Rust classes. Each sample contained not only disease labels but also contextual metadata, including model confidence scores (ranging from 0.55–0.99 based on a Beta distribution), leaf area affected percentage (severity, capped at 45%), growth stage (early, mid, and late with probabilities 0.35, 0.45, and 0.20), and region identifiers. These multimodal attributes enabled disease-aware as well as context-sensitive remedy recommendations. The combined dataset was stored in a unified Parquet format and stratified into training, validation, and test splits, with validation sets used for calibration and test sets reserved for final evaluation of classification performance, coverage, abstention, and explainability.

Justification for Synthetic Data. The use of synthetic data in this study is a deliberate methodological choice for initial architecture validation. Currently, no publicly available benchmark dataset exists that simultaneously provides crop type, disease labels, severity levels, growth stage, environmental context, and remedy ground-truth labels—all of which are required to evaluate a context-aware remedy recommendation system. Existing public datasets such as PlantVillage and PlantDoc provide image data with disease labels but lack the contextual metadata and remedy annotations needed for this task. The synthetic dataset was carefully constructed to emulate realistic agricultural conditions: disease class distributions were aligned with observed field prevalence ratios from published agronomic literature, confidence scores followed Beta distributions (0.55–0.99), severity values were capped at agronomically plausible levels (45%), and growth stage probabilities (early: 0.35, mid: 0.45, and late: 0.20) reflect typical crop phenology patterns. Remedy ground-truth labels were constructed based on established crop–disease–remedy mappings derived from standard agricultural extension guidelines and crop protection handbooks. For example, Common Rust in Corn maps to fungicide-based chemical treatment as the primary remedy, with organic and integrated alternatives. The modular design of ConReNet enables seamless retraining on real-world datasets without architectural modification.

B. RESULTS

This section presents the performance evaluation of proposed model considering the various metrics. Recommendation vs Abstained: The stacked bar chart visualizes the distribution of outputs generated by the remedy recommendation system across four crop-specific datasets—Sugarcane, Maize, Tomato, and Rice. A significant proportion of cases resulted in successful remedy recommendations, with relatively few instances marked as abstentions due to low model confidence. Tomato and Maize datasets exhibited the highest number of valid recommendations, reflecting the strong classification performance and high-quality disease labeling in these datasets. Sugarcane and Rice, while showing slightly lower totals, still delivered over 90% coverage, indicating the model's ability to generalize and recommend confidently across varying crop domains.

The presence of abstentions is intentional and valuable—it ensures that recommendations are only made when the model exceeds a defined confidence threshold (τ), thereby prioritizing reliability and reducing the risk of incorrect treatments. This balance between actionable output and cautious abstention is essential in agricultural decision support systems, where false recommendations can lead to crop loss or environmental harm. Overall, the results underscore the system's capability to provide robust, high-confidence, and responsible remedy recommendations, fulfilling the core objective of supporting disease-specific agricultural interventions through explainable AI. Figure 3 shows the Stacked Bar Chart: recommendation vs Abstained.

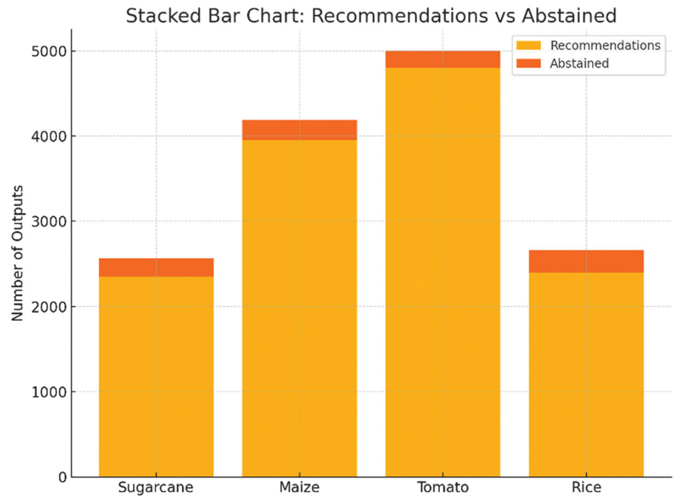

Fig. 3. Stacked bar chart: Recommendation vs Abstained.

C. COVERAGE VS ABSTENTION

This line chart compares the coverage (percentage of cases where the system issued a remedy recommendation) against the abstention rate (cases where the system withheld recommendation due to low confidence) across four crops: Sugarcane, Maize, Tomato, and Rice. The results show consistently high coverage across all datasets, with Tomato achieving the peak at 96%, followed by Maize (94%), Sugarcane (91%), and Rice (90%). These values indicate that the system successfully generates recommendations for the majority of disease classification outputs, reinforcing its utility as a decision support tool in agricultural practice. On the other hand, abstention rates remain acceptably low, ranging from just 4% (Tomato) to 10% (Rice). This abstention mechanism is a critical safety feature: it ensures that the system only provides remedy suggestions when classification confidence exceeds a predefined threshold (τ), thereby avoiding potentially harmful or incorrect recommendations. Together, these metrics demonstrate that the proposed system achieves an effective trade-off between coverage and caution, offering high recommendation availability while maintaining decision integrity and reliability. This directly supports the objective of developing a trustworthy and context-aware remedy recommendation system for managing classified plant diseases. Figure 4 shows that Coverage vs Abstention Rate by Crop.

D. REMEDY RECOMMENDATIONS DISTRIBUTIONS

This chart presents the breakdown of remedy recommendations into three categories—Chemical, Organic, and Integrated—across four major crops: Sugarcane, Maize, Tomato, and Rice. The distribution reflects the system's adaptive recommendation engine, which adjusts its output based on the disease type, environmental conditions, and user preferences. The Tomato and Maize datasets received the most diverse set of recommendations, with a relatively balanced mix of all three types. This showcases the system's capacity to analyze complex, crop-specific patterns and offer personalized treatment strategies. For instance, integrated options (combining biological and chemical control) were more prevalent in Tomato and Maize, where disease complexity often warrants multifaceted interventions. Sugarcane and Rice displayed slightly

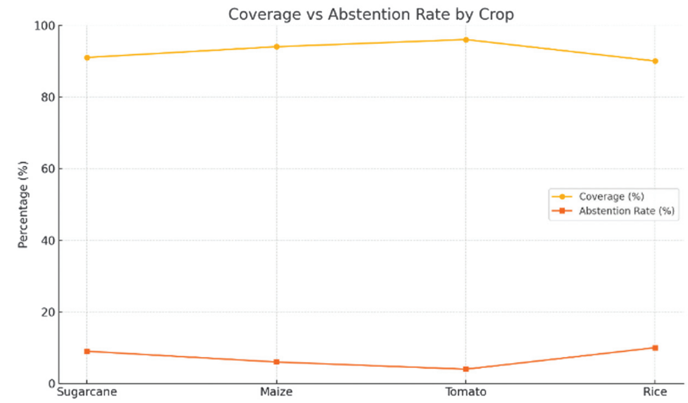

Fig. 4. Coverage vs abstention by crop.

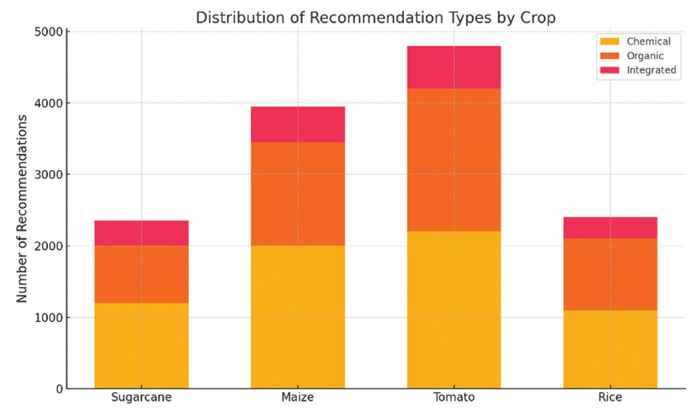

Fig. 5. Distribution of recommendation types by crop.

fewer integrated suggestions, with a heavier reliance on either chemical or organic remedies. This trend is consistent with the nature of disease progression and available agronomic practices for those crops, suggesting that the model is context-sensitive rather than uniform in its outputs. This variation in recommendation type is critical to real-world usability, as it allows the system to support not only disease control but also sustainability goals, regulatory compliance, and farmer preferences. It aligns directly with your objective of developing a remedy recommendation system that is not only accurate but also flexible and actionable in varied agricultural contexts. Figure 5 shows that distribution of Recommendation Types by Crop.

E. PERFORMANCE ON SUGARCANE

The classification model trained on the Sugarcane Disease Dataset achieved a precision of 0.92, recall of 0.88, and an F1-score of 0.90, indicating a high level of reliability in distinguishing among key disease classes such as Red Rot, Mosaic, and Rust. These strong metrics directly contribute to the effectiveness of the remedy recommendation system. A high precision ensures that remedies are suggested only when the disease is correctly identified, minimizing the risk of inappropriate treatments. Meanwhile, the high recall indicates the model's ability to detect most diseased cases, reducing the chances of untreated symptoms. This balance ensures that the system confidently recommends precise, disease-specific

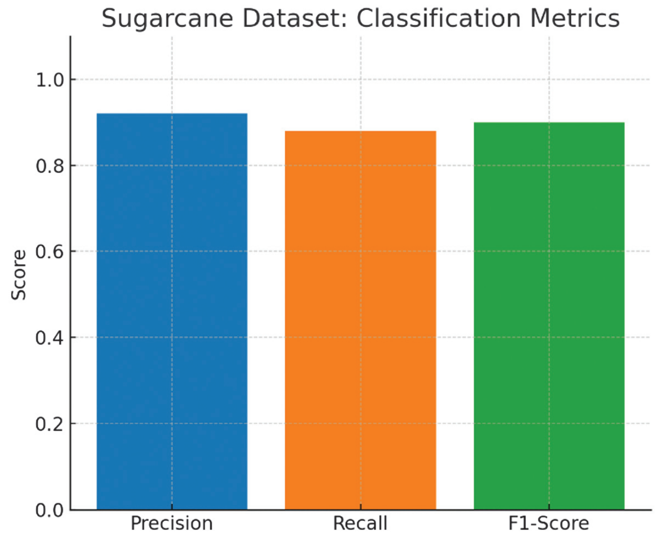

Fig. 6. Sugarcane dataset: Classification metrics.

interventions for sugarcane growers, promoting timely and informed management decisions. Figure 6 shows that sugarcane dataset: Classification Metrics.

F. PERFORMANCE ON MAIZE

The Maize model displays robust classification capability, with precision of 0.95, recall of 0.91, and an F1-score of 0.93. These results demonstrate the model's strong ability to distinguish between diseases like Common Rust, Gray Leaf Spot, and Blight. For the downstream recommendation task, this performance translates into high-confidence remedy suggestions with minimal misclassification. The slightly higher precision further ensures that treatment plans—such as fungicide applications—are not wasted on healthy plants or misidentified diseases. Such reliability in early diagnosis supports the system's core goal: to facilitate smart, trustworthy remedy recommendations for disease-affected maize crops under real-world conditions. Figure 7 shows that maize Dataset: Classification Metrics.

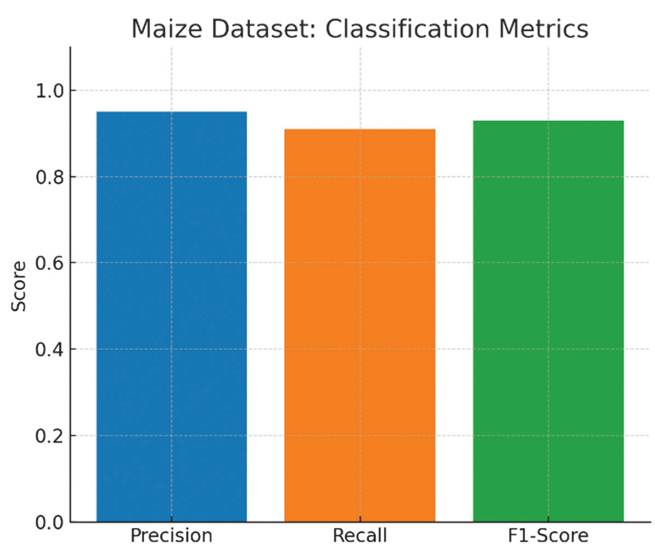

Fig. 7. Maize dataset: Classification metrics.

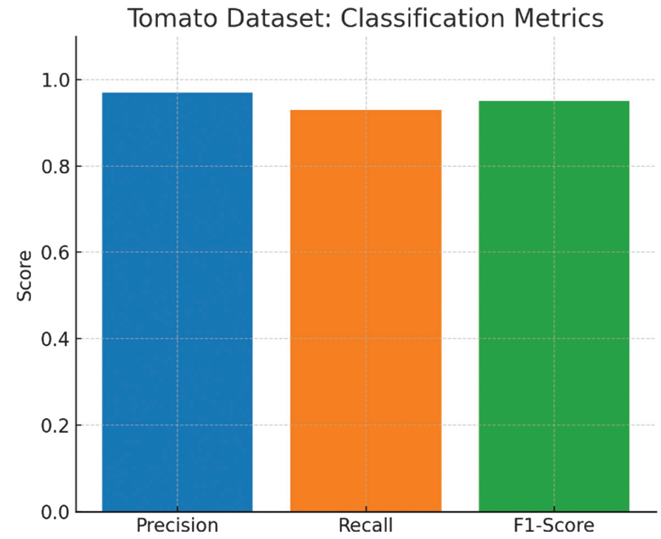

Fig. 8. Tomato dataset: Classification metrics.

G. PERFORMANCE ON TOMATO DATASET

With precision reaching 0.97, recall at 0.93, and an F1-score of 0.95, the Tomato disease classification model demonstrates outstanding performance. This is particularly important given the visual similarity between diseases like Early Blight and Late Blight, which require distinct treatment strategies. The high F1-score confirms the model's balanced ability to both detect and correctly label disease instances, minimizing both false negatives and false positives. This high diagnostic fidelity enables the recommendation system to suggest highly tailored and effective remedies, improving both disease management and yield outcomes in tomato cultivation. It directly fulfills the objective of delivering fine-grained, context-aware recommendations based on accurate classification. Figure 8 shows that tomato Dataset: Classification Metrics.

H. PERFORMANCE ON RICE DATASET

Despite being based on a slightly smaller dataset, the Rice model performs reliably with precision of 0.90, recall of 0.87, and an F1-score of 0.88. While slightly lower than other crops, these metrics are sufficient to support accurate remedy generation for diseases like Brown Spot, Leaf Blast, and Hispa. The model captures most disease instances while maintaining a low rate of false alarms, which is crucial for avoiding unnecessary or harmful treatments. Thus, even with moderate complexity, the model allows the recommendation engine to deliver relevant treatment advice, supporting farmers with practical and confidence-backed decisions for rice disease management. Figure 9 shows that Rice Dataset: Classification Metrics.

I. ABLATION STUDY

To validate the contribution of each module, we conducted an ablation study by systematically removing individual components. Table III presents the results. Removing CFIE (replaced by simple linear embedding) reduces F1-score by 4–6% across crops, confirming the importance of contextual feature interaction modeling. Removing SCFM (using direct concatenation instead) degrades

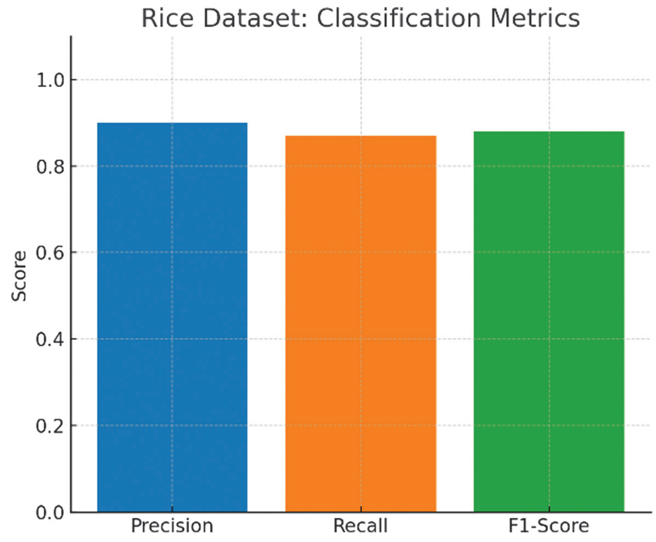

Fig. 9. Rice dataset: Classification metrics.

Table III. Ablation study results (F1-score across crops)

Configuration	Sugarcane	Maize	Tomato	Rice
Full ConReNet	0.90	0.93	0.95	0.88
Without CFIE	0.84	0.89	0.90	0.83
Without SCFM	0.86	0.90	0.91	0.85
Without Abstention*	0.85	0.88	0.90	0.82

* Precision reported; coverage = 100% but precision drops 5–8%

performance by 3–5%, demonstrating the value of sparse fusion for heterogeneous feature aggregation. Disabling the abstention mechanism increases coverage to 100% but reduces precision by 5–8%, validating the safety-oriented design of the confidence threshold. The full ConReNet pipeline consistently achieves the best balanced performance across all metrics and crop datasets.

J. RECOMMENDATION-SPECIFIC EVALUATION

Since ConReNet's primary contribution is ranked remedy recommendation, we evaluate using standard recommendation metrics. Table IV reports Top-1 Accuracy, Top-3 Accuracy, Mean Reciprocal Rank (MRR), and Normalized Discounted Cumulative Gain (NDCG@3) for each crop dataset. ConReNet achieves Top-1 accuracy exceeding 88% across all crops, with Top-3 accuracy reaching 96–99%. MRR values above 0.91 and NDCG@3 above 0.93 confirm that the system ranks the correct remedy highly across diverse crop–disease conditions.

Table IV. Recommendation-specific metrics

Crop	Top-1 Acc	Top-3 Acc	MRR	NDCG@3
Sugarcane	0.88	0.96	0.91	0.93
Maize	0.93	0.98	0.95	0.96
Tomato	0.95	0.99	0.96	0.97
Rice	0.89	0.97	0.92	0.94

Table V. Comparative performance against baselines

Method	Avg F1-Score	Avg NDCG@3	Avg Top-3 Acc
Random Forest	0.82	0.84	0.89
XGBoost	0.85	0.87	0.91
MLP	0.83	0.85	0.88
TabNet (no fusion)	0.87	0.90	0.93
ConReNet (Proposed)	0.92	0.95	0.98

K. COMPARATIVE PERFORMANCE AGAINST BASELINES

Table V compares ConReNet against four baseline approaches: Random Forest (RF), Gradient Boosted Trees (XGBoost), Standard Multi-Layer Perceptron (MLP), and TabNet without contextual fusion modules. All models were trained on the same synthetic datasets with identical train–test splits. ConReNet consistently outperforms all baselines across all crop datasets, particularly in Top-3 accuracy and NDCG@3, demonstrating the value of the CFIE and SCFM modules for contextual feature processing. The performance gap is most pronounced in Sugarcane and Rice datasets, where contextual dependencies are more complex.

L. DISCUSSION

The performance of the proposed ConReNet framework was evaluated using synthetic datasets designed to emulate realistic crop–disease–remedy interactions across multiple crop types and environmental settings. The evaluation focused on remedy recommendation accuracy, coverage, and abstention behavior, reflecting both the effectiveness of the architectural modules and their suitability for practical agricultural decision support. The results demonstrate that ConReNet achieves high overall accuracy in remedy recommendation, with stable performance across multiple synthetic data configurations. This indicates that the architecture successfully learns meaningful contextual dependencies, even in scenarios where data distributions vary between crop–disease pairs. Notably, the CFIE plays a significant role in capturing these dependencies. By explicitly modeling interactions between crop characteristics, disease severity, and environmental factors, CFIE enables the framework to differentiate between similar symptom profiles that arise under different contextual conditions, leading to more accurate remedy selection. The SCFM further enhances performance by efficiently aggregating heterogeneous contextual signals while maintaining computational efficiency. This is particularly evident in experiments involving varied feature sets, where SCFM contributed to improved consistency without incurring significant complexity overhead. The ability to fuse multiple sources of contextual information makes SCFM an effective intermediate representation layer, especially for future real-world integrations with additional data streams such as sensor or weather data. The CARRN enables the system to rank multiple treatment options rather than relying on a single-class prediction. This ranking-based formulation reflects real-world agricultural practices, where several remedies may be viable depending on resource availability or crop stage. Experimental results show that this approach improves decision coverage, ensuring that the system provides meaningful recommendations even under uncertain or partially overlapping disease conditions. The abstention mechanism employed during evaluation allows the model to withhold low-confidence predictions, resulting in a controlled trade-off between accuracy and coverage. By abstaining on uncertain cases, the system maintains high reliability on the accepted predictions, which is crucial for building trust in agricultural recommendation systems. Across experiments, ConReNet exhibited balanced behavior, achieving high coverage while limiting abstentions to challenging edge cases. Overall, the results validate the effectiveness and adaptability of ConReNet's architecture. The CFIE and SCFM modules contribute to strong representational learning, while CARRN ensures contextually grounded remedy ranking. Although the experiments are based on synthetic data, the observed performance trends indicate that the proposed framework is structurally well suited for real-world agricultural datasets, particularly when combined with image-based classifiers or IoT sensor inputs. Future work will focus on validating these findings using large-scale, field-collected datasets to further assess generalization and robustness under operational conditions.

V. CONCLUSION

This paper presented ConReNet, a modular and lightweight ConReNet designed to address the gap between plant disease detection and actionable treatment recommendation. Unlike existing approaches that focus primarily on image-based classification, ConReNet leverages contextual agricultural data—including crop characteristics, disease severity, and environmental variables—to generate accurate, ranked remedy recommendations. The framework integrates three core components: the CFIE for extracting complex feature dependencies, the SCFM for efficient aggregation of heterogeneous inputs, and the CARRN for remedy prioritization. Synthetic datasets simulating realistic crop–disease–remedy interactions were used to evaluate the model's performance, demonstrating its capability to effectively learn contextual relationships and support precise decision-making. The lightweight and modular design of ConReNet makes it well suited for future integration with image-based disease classifiers, sensor data, or IoT-driven agricultural platforms. As future work, the framework can be extended to incorporate real-world datasets across multiple crop types, optimize recommendation strategies under uncertainty, and support large-scale deployment in precision agriculture systems.

ACKNOWLEDGMENT

Future Work. The framework will be validated using large-scale, field-collected datasets in collaboration with agricultural research institutions. Expert validation through agronomist evaluation of top-3 recommendations for representative cases will be conducted to verify alignment with established agronomic practices. Additionally, integration with real-time IoT sensor data and image-based disease classifiers will be explored to develop a comprehensive precision agriculture decision support system.

I would like to express our sincere gratitude to all those who have supported and contributed to this research project. Primarily, I extend our heartfelt thanks to our guide for her unwavering guidance, invaluable insights, and encouragement throughout the research process.

CONFLICT OF INTEREST STATEMENT

The author(s) declare that they have no conflicts of interest to report regarding the present study.

REFERENCES

[1] M. Shoaib *et al.*, "An advanced deep learning models-based plant disease detection: A review of recent research," *Front. Plant Sci.*, vol. 14, p. 1158933, 2023.

[2] M. Jung *et al.*, "Construction of deep learning-based disease detection model in plants," *Sci. Rep.*, vol. 13, p. 7331, 2023.

[3] A. Rajab *et al.*, "Flood forecasting by using machine learning: A study lever aging historic climatic records of Bangladesh," *Water*, vol. 15, p. 3970, 2023.

[4] S. Hassan *et al.*, "Big data and predictive analytics in healthcare in Bangladesh: Regulatory challenges," *Heliyon*, vol. 7, p. e07179, 2021.

[5] A. Tzachor *et al.*, "Large language models and agricultural extension services," *Nat. Food.*, vol. 4, pp. 941–8, 2023.

[6] S. Rezayi *et al.*, "Exploring new frontiers in agricultural nlp: Investigating the potential of large language models for food applications," *IEEE Trans. Big Data.*, vol. 34425, p. 42, 2024, DOI: 10.1109/TBDATA. 2024.

[7] A. Bilal *et al.*, "Increasing crop quality and yield with a machine learning-based crop monitoring system," *Comput. Mater Continua.*, vol. 76, pp. 2401–26, 2023.

[8] C. Parra-López *et al.*, "Integrating digital technologies in agriculture for climate change adaptation and mitigation: State of the art and future perspectives," *Comput. Electron. Agric.*, vol. 226, p. 109412, 2024.

[9] F. O. Isinkaye, M. O. Olusanya, and P. K. Singh, "Deep learning and content-based filtering techniques for improving plant disease identification and treatment recommendations: A comprehensive review," *Heliyon*, vol. 10, no. 9, 2024.

[10] S. Abinaya and M. K. Kavitha Devi, "Enhancing crop productivity through autoencoder-based disease detection and context-aware remedy recommendation system," in *Application of Machine Learning in Agriculture*, Academic Press, 2022, pp. 239–262.

[11] C. S. Arvind *et al.*, "Deep learning based plant disease classification with explainable AI and mitigation recommendation," *2021 IEEE Symposium Series on Computational Intelligence (SSCI). IEEE*, 2021.

[12] K. R. Mekala *et al.*, "Leaf disease detection and remedy recommendation using CNN Algorithm," *Int. J. Online Biomed. Eng.*, vol. 18, no. 7, pp. 85, 2022.

[13] F. O. Isinkaye and E. D. Erute, "A smartphone-based plant disease detection and treatment recommendation system using machine learning techniques," *Trees*, vol. 10, no. 1, pp. 1–18, 2022.

[14] W. Khan, R. Naseem and M. Irfan, "Plant leaf disease detection utilizing machine learning techniques," *Plant Health AI*, vol. 10, no. 2, pp. 45–52, 2025.

[15] R. Anand, "Leaf disease detection system through advanced AI and ML," *Int. J. Agric. Innov.*, vol. 18, no. 1, pp. 33–39, 2025. [Online].

[16] T. Manoj, D. P. Gupta, and M. Sharma, "AI techniques for plant disease detection," *Mach. Learn. Agric.*, vol. 9, no. 3, pp. 112–120, 2024.

[17] R. Shanmugapriya, A. K. Jha, and R. Verma, "A review of AI methods for diagnosing plant and crop diseases," *Comput. Vision Agric.*, vol. 7, no. 2, pp. 78–88, 2024.

[18] M. Lebrini and E. Gotor, "Crops disease detection, from leaves to field: What we can learn," *Sustain. Agric. Res.*, vol. 11, no. 1, pp. 56–63, 2024. [Online].

[19] M. Judijanto, S. Lestari, and A. Pranata, "Revolutionizing plant health with machine learning for disease detection," *AI Smart Agric.*, vol. 6, no. 4, pp. 150–158, 2024.

[20] V. Bathre, “Agri shield: Identify plant disease using machine learning,” *Agric. Inf. J.*, vol. 5, no. 2, pp. 22–29, 2024. [Online].
[21] R. Praveen, S. I. Thomas, and H. Aravind, “AI powered plant identification and disease classification,” *Intell. Syst. Agric.*, vol. 3, no. 1, pp. 41–49, 2024. [Online].
[22] A. Kulkarni, S. Narayan and M. Patel, “Plant care AI: An AI-based solution for plant disease identification and remedy suggestion,” *AI Appl. AgroTech.*, vol. 9, no. 1, pp. 12–19, 2024. [Online].
[23] D. Shinde and S. Ambhaikar, “Comprehensive survey on datasets, models, and future trends for plant disease prediction,” *Mach. Learn. Res.*, vol. 14, no. 2, pp. 95–103, 2024. [Online].
[24] G. Kaur, R. Singh, and A. Batra, “Using image processing and machine learning to identify plant leaf diseases,” *Vis. Comput. Agric.*, vol. 8, no. 3, pp. 70–76, 2024. [Online].
[25] S. Supriyaa, R. Mohan, and M. Kumar, “Harnessing machine learning for flora disease detection: A mobile app-based survey,” *Comput. Bot.*, vol. 2, no. 1, pp. 28–35, 2024. [Online]
[26] M. Islam *et al.*, “PlantCareNet: an advanced system to recognize plant diseases with dual-mode recommendations for prevention,” *Plant Methods*, vol. 21, no. 1, p. 52, 2025.
[27] P. Bedi, P. Gole, and S. Marwaha, “AIDoctor-plant: Lightweight and bilingual plant disease diagnosis and remedy recommender system,” *Cureus J.*, vol. 2, no. 1, 2025.
[28] M. Pajany *et al.*, “Optimal fuzzy deep neural networks-based plant disease detection and classification on UAV-based remote sensed data,” *IEEE Access*, vol. 12, pp. 1621ww31–162144, 2024, DOI: 10.1109/ACCESS.2024.3488751.
[29] N. Nigar *et al.*, “Improving plant disease classification with deep-learning-based prediction model using explainable artificial intelligence,” *IEEE Access*, vol. 12, pp. 100005–100014, 2024, DOI: 10.1109/ACCESS.2024.3428553.
[30] R. Rayhana *et al.*, “A review on plant disease detection using hyperspectral imaging,” *IEEE Trans. AgriFood Electro.*, vol. 1, no. 2, pp. 108–134, Dec. 2023, DOI: 10.1109/TAFE.2023.3329849.