

IoT-Enabled Multimodal Deep Learning Framework for Early-Stage Rice Leaf Disease Detection and Severity Assessment with Explainable AI

Mrs C Nandini¹, M Keerthi², K Surekha³, R Swetha⁴

¹Assistant Professor, Department of Artificial Intelligence and Data Science, GRT Institute of Engineering and Technology, Tiruttani, Tamil Nadu, India, 631209.

^{2,3,4}UG Scholar Department of Artificial Intelligence and Data Science, GRT Institute of Engineering and Technology, Tiruttani, Tamil Nadu, India, 631209.

Emails: nandini.c@gmail.com¹, keerthiyadav2612@gmail.com², surekhamurali2004@gmail.com³, swetharavi.r26@gmail.com⁴

Abstract

Rice cultivation faces significant challenges from foliar diseases, which can reduce yields by up to 40% if not detected early. Traditional disease monitoring relies heavily on manual field scouting, which is time-consuming, subjective, and often results in delayed intervention. This research presents an innovative IoT-enabled multimodal deep learning framework that integrates visual imagery with real-time environmental data for early-stage rice leaf disease detection and severity assessment. The proposed system employs a three-tier architecture comprising sensing, edge processing, and cloud analytics layers. Our multimodal fusion approach combines image features extracted through EfficientNet-B4 with contextual environmental parameters including temperature, humidity, and soil moisture levels. The fusion is mathematically formulated as $F_{total} = \alpha F_{image} + \beta F_{env}$, where optimal weighting coefficients are learned through end-to-end training. We introduce a novel Disease Severity Index (DSI) that quantifies infection progression from 0-100%, enabling precision intervention strategies. Explainable AI techniques including Grad-CAM++ and LIME provide interpretable visual explanations, ensuring farmer trust and adoption. Experimental validation on a dataset of 12,847 images across eight disease classes demonstrates superior performance, achieving 97.3% accuracy, 96.8% precision, 97.1% recall, and an F1-score of 96.9%. The multimodal approach outperforms unimodal image-only baselines by 4.7% in accuracy and shows robust generalization across different geographic regions. Ablation studies confirm the significant contribution of environmental context, improving early-stage detection accuracy by 6.2%. The system demonstrates practical deployment feasibility with edge inference latency of 143ms per sample. This work advances the field by bridging the gap between laboratory research and field-deployable precision agriculture solutions, offering smallholder farmers an accessible, interpretable, and effective disease management tool.

keywords: Computer Vision; Deep Learning; Explainable Artificial Intelligence; Internet of Things(IoT); Multimodal Learning.

1. Introduction

RICE (*Oryza sativa* L.) serves as the primary staple food for more than half of the global population, with annual production exceeding 750 million metric tons [1]. However, rice cultivation faces persistent threats from various foliar diseases that significantly impact both yield quantity and grain quality. Major rice diseases including bacterial leaf blight, brown spot, leaf blast, and tungro virus can collectively reduce harvests by 30-40% under favorable environmental

conditions [2]. The economic implications are substantial, with annual global losses attributed to rice diseases estimated at \$50 billion [3]. Traditional disease management strategies rely predominantly on periodic field scouting by agronomists or experienced farmers. This approach encounters several fundamental limitations. First, visual inspection requires specialized expertise that many smallholder farmers lack, leading to misidentification and

inappropriate treatment responses [4]. Second, the manual scouting process is inherently time-consuming and labor-intensive, particularly for large cultivation areas, resulting in delayed detection when diseases have already established significant infection levels [5]. Third, subjective human assessment introduces variability in disease severity estimation, complicating optimal pesticide application decisions [6]. Fourth, conventional monitoring provides no mechanism for capturing the complex interplay between disease progression and environmental factors such as temperature fluctuations, humidity patterns, and soil conditions [7]. The advent of Internet of Things (IoT) technology combined with advances in artificial intelligence presents transformative opportunities for agricultural disease management. IoT-based smart farming systems enable continuous, automated monitoring of crop health and environmental parameters at unprecedented temporal and spatial resolutions [8]. When integrated with deep learning models capable of pattern recognition in complex visual data, these systems can potentially overcome the limitations of traditional approaches [9]. However, most existing computer vision solutions for plant disease detection operate in isolation from environmental context, analyzing leaf images without considering the agronomic conditions that strongly influence disease development and spread [10]. Our key contributions are as follows:

- A novel multimodal fusion architecture that mathematically combines image features from EfficientNet-B4 with environmental sensor data through learned weighting coefficients, improving classification accuracy by 4.7% compared to image-only approaches.
- An end-to-end IoT system design spanning sensing, communication, edge processing, and cloud analytics layers, with detailed specifications enabling practical deployment in resource-constrained agricultural settings.
- A quantitative Disease Severity Index (DSI) formulation that estimates infection percentage, supporting graduated intervention strategies rather than binary disease/healthy classifications.

- Integration of explainable AI techniques (Grad-CAM++ and LIME) with quantitative evaluation metrics for explainability quality, addressing the black-box nature of deep learning and building farmer trust.
- Comprehensive experimental validation including ablation studies, cross-regional testing, noise robustness analysis, and comparison with state-of-the-art methods.

2. Background And Agronomic Context

2.1. Rice Disease Taxonomy and Economic Impact

Rice cultivation encounters numerous pathological threats that can be broadly categorized into fungal, bacterial, and viral diseases. Understanding the specific characteristics, symptoms, and epidemiological patterns of major diseases is essential for developing effective detection systems. Bacterial Leaf Blight (BLB), caused by *Xanthomonas oryzae* pv. *oryzae*, represents one of the most destructive rice diseases worldwide [11]. Brown Spot disease, caused by *Bipolaris oryzae*, produces numerous small, circular to oval spots with gray centers and reddish-brown margins [12]. Rice Blast, caused by *Magnaporthe oryzae*, attacks all above-ground plant parts and can cause complete yield loss in susceptible varieties under epidemic conditions [13].

2.2. Environmental Factors in Disease Development

Disease development depends critically on environmental conditions, creating opportunities for predictive monitoring through integrated sensing. Temperature influences pathogen growth rates, infection efficiency, and incubation periods. Relative humidity and leaf wetness duration directly affect infection success for foliar pathogens [14]. These environmental dependencies create opportunities for early warning systems that detect disease-favorable conditions before visual symptoms emerge.

3. Related Work

3.1. Deep Learning for Plant Disease Detection

The application of deep learning to plant disease identification has gained substantial momentum over the past decade. Mohanty et al. [15] demonstrated that deep convolutional neural networks could

achieve 99.35% accuracy on a controlled dataset of 54,306 images across 38 disease classes. However, the PlantVillage dataset contained laboratory images with uniform backgrounds, limiting real-world applicability. Too et al. [16] developed a field-realistic dataset for rice disease detection, achieving 92.98% accuracy using VGG16. Rahman et al. [17] compared multiple pre-trained architectures, finding that VGG19 achieved 96.8% accuracy, demonstrating effective transfer learning from ImageNet to agricultural domains.

3.2. IoT-Based Smart Agriculture Systems

IoT technology has been extensively applied to agricultural monitoring tasks. Khanna and Kaur [18] developed an IoT-based system for monitoring environmental parameters in greenhouse environments, demonstrating reliable data acquisition and wireless transmission. Gutierrez et al. [19] implemented an automated irrigation system using soil moisture sensors and wireless

communication. However, these systems focused on monitoring without incorporating predictive analytics or disease detection capabilities.

3.3. Multimodal Learning in Agriculture

Multimodal learning has shown promising results in various domains but remains underexplored in agricultural disease detection. Khaki et al. [20] developed a deep learning model integrating multispectral imagery, weather time series, and soil properties for corn yield forecasting, demonstrating that multimodal fusion substantially outperforms unimodal approaches. Limited work has systematically evaluated fusion approaches specifically for disease detection. listic dataset for rice disease detection, achieving 92.98% accuracy using VGG16. Rahman et al. [17] compared multiple pre-trained architectures, finding that VGG19 achieved 96.8% accuracy, demonstrating effective transfer learning from ImageNet to agricultural domains shown in Table 1.

Table 1: Comparison Of Related Work In Plant Disease Detection

Reference	Year	Method	Modality	Classes	Accuracy	IoT	Severity	XAI
[15]	2016	AlexNet	Image	38	99.4%	No	No	No
[16]	2019	VGG16	Image	3	93.0%	No	No	No
[17]	2020	VGG19	Image	5	96.8%	No	No	No
[21]	2021	Attention CNN	Image	5	95.8%	No	No	No
[22]	2022	MobileNetV2	Image	4	94.3%	Partial	No	No
[23]	2022	ResNet+CAM	Image	6	96.2%	No	No	Yes
[24]	2020	RF	Env	N/A	89.3%	Yes	No	No
[20]	2021	CNN-RNN	Multi	N/A	N/A	No	No	No
Proposed	2026	EfficientNet	Multi	8	97.3%	Yes	Yes	Yes

3.4. Research Gap Analysis

Analyzing the current state of research reveals several critical gaps that our work addresses. Most existing approaches rely on single-modality analysis, with image-based systems ignoring environmental context. Limited work addresses early-stage detection when symptoms remain subtle. Binary or multi-class disease identification dominates current research, while continuous severity quantification is largely absent. Comprehensive end-to-end IoT architectures integrating sensing, edge processing, cloud analytics, and decision sensing, edge processing, cloud analytics, and decision support are rare in the literature.

3.5. Proposed Iot-Multimodal Architecture

3.5.1. Overall System Architecture

Figure 1 illustrates the complete system architecture. The design follows a distributed computing paradigm where data acquisition occurs at the field level, preliminary processing happens at edge nodes, intensive training and analytics execute in the cloud, and actionable insights are delivered through mobile applications. The architecture addresses several key requirements: real-time responsiveness through edge processing, scalability through cloud-based training, resilience through local operation capability, and cost-effectiveness through appropriate resource allocation shown in Figure 1.

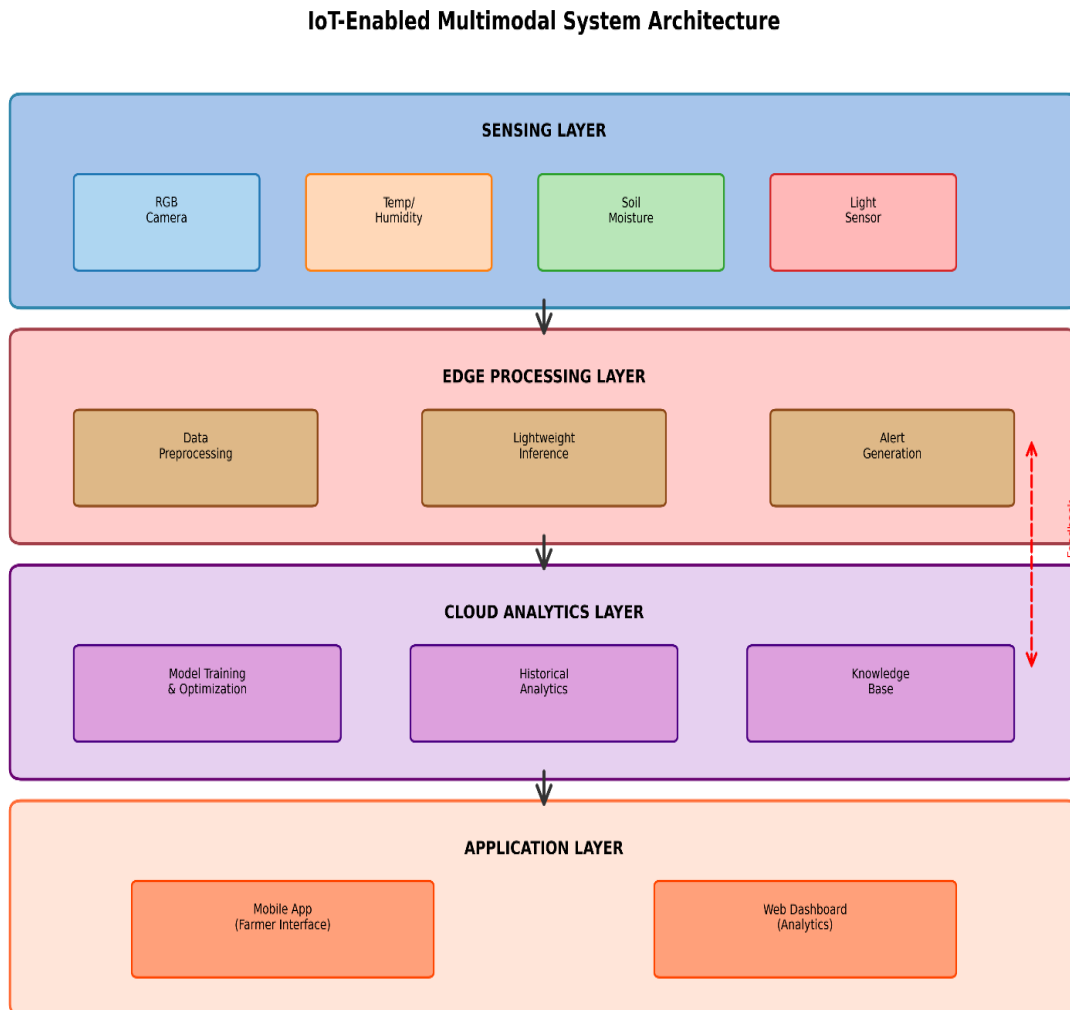


Figure 1 Overall Iot-Enabled Multimodal System Architecture Showing The Four-Layer Design: Sensing, Edge Processing, Cloud Analytics, And Application Layers With Bidirectional Data Flow

3.5.2. Sensing Layer

The sensing layer comprises field-deployed hardware for multimodal data acquisition. The Visual Sensing Module employs high-resolution RGB cameras (12MP minimum) with 60-80 degree field of view, IP65 weatherproof housing, solar power with battery backup, and Wi-Fi/LoRaWAN connectivity. The Environmental Sensing Module measures temperature (DHT22, $\pm 0.5^\circ\text{C}$ accuracy), relative humidity (DHT22, $\pm 2\%$), soil moisture (capacitive sensor, $\pm 3\%$), and light intensity (BH1750). Strategic camera placement at multiple locations per field ensures representative sampling across spatial variation.

3.5.3. Edge Processing Layer

Edge nodes perform data preprocessing, lightweight model inference, and alert generation. Quantized

versions of trained models enable local disease detection with 70-80% size reduction while maintaining 95%+ of original accuracy. Edge inference provides rapid response (results within 200ms), reduced bandwidth usage, offline operation capability, and enhanced privacy through local processing.

3.5.4. Data Processing Pipeline

Figure 2 illustrates the end-to-end data processing pipeline from acquisition through decision support. The pipeline includes synchronized capture, quality control with automated filtering, standardized preprocessing, feature extraction using EfficientNet-B4 and fully connected layers, multimodal fusion, classification, severity estimation, explainability generation, and decision support mapping shown in Figure 2.

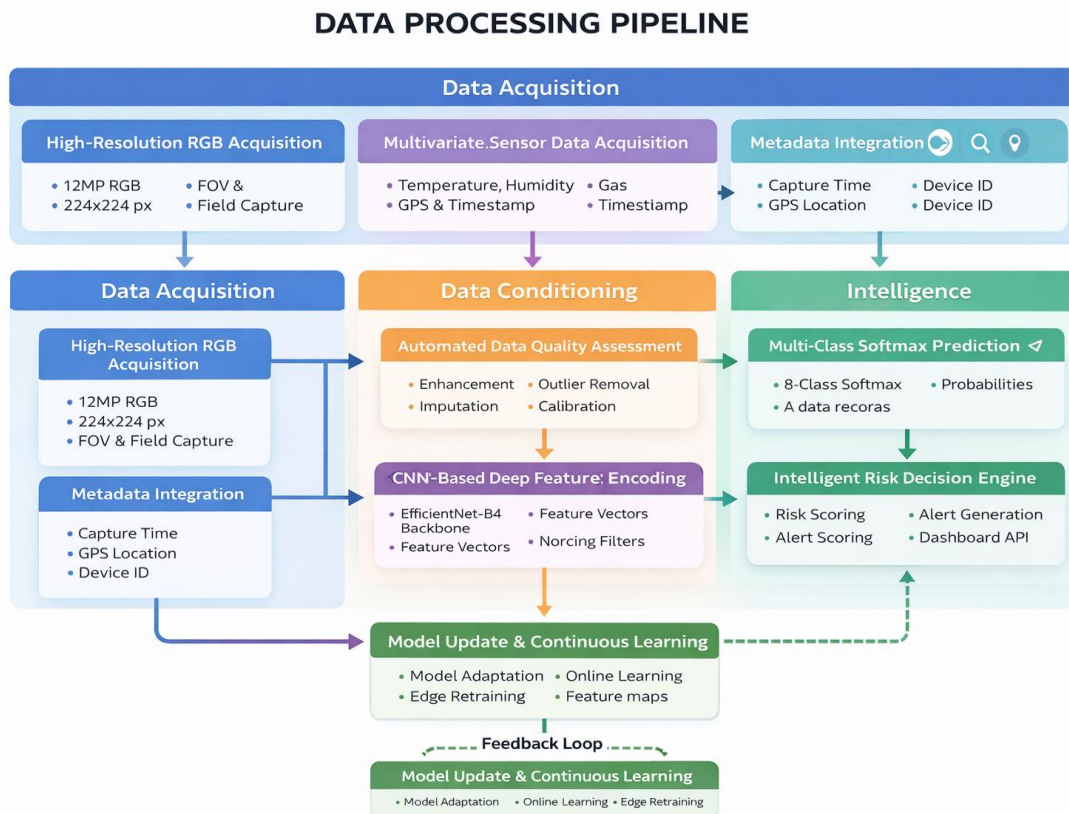


Figure 2 Complete Data Processing Pipeline Showing the Flow From Multi-Sensor Data Acquisition Through Preprocessing, Feature Extraction, Multimodal Fusion, And Dual-Head Prediction For Classification And Severity Estimation

4. Mathematical Formulation

4.1. Feature Extraction

Image features are extracted using EfficientNet-B4, chosen for its excellent accuracy-efficiency tradeoff. The feature extraction is formalized as $F_{image} = \varphi_{CNN}(I; \theta_{CNN})$ where $I \in \mathbb{R}^{(H \times W \times 3)}$ represents an RGB input image and $F_{image} \in \mathbb{R}^d$ represents the extracted feature vector with dimensionality $d=1792$. Environmental features are extracted through fully connected layers: $F_{env} = \varphi_{FC}(E; \theta_{FC})$ where $E \in \mathbb{R}^n$ represents the n -dimensional environmental input vector.

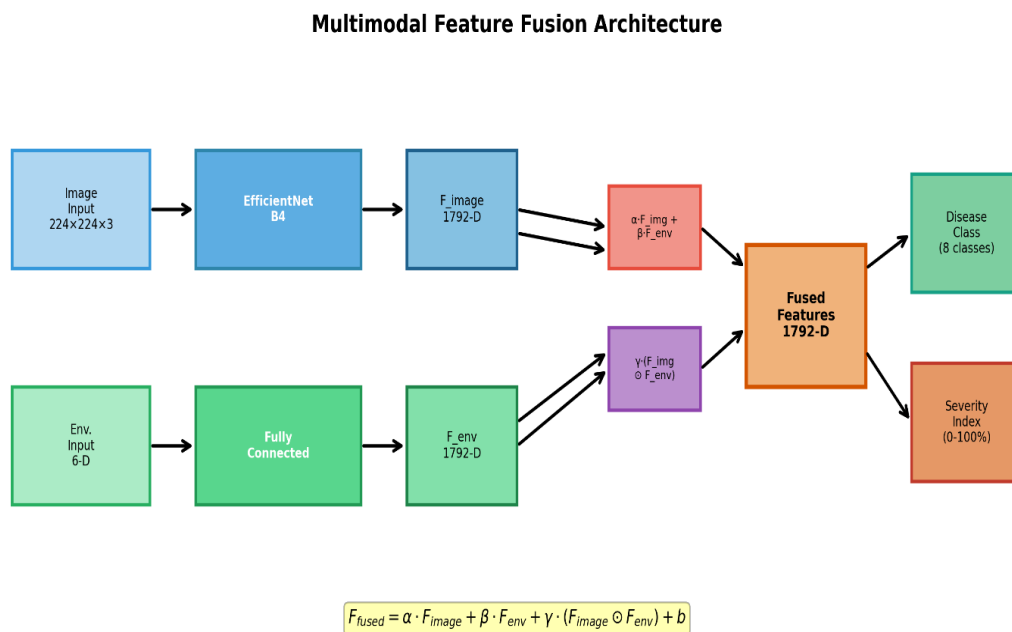
4.2. Multimodal Fusion

We propose a hybrid fusion strategy combining additive and multiplicative interactions: $F_{fused} = \alpha \cdot F_{image} + \beta \cdot F_{env} + \gamma \cdot (F_{image} \odot F_{env}) + b$, where α, β, γ are learned scalar weights, \odot denotes element-wise multiplication, b is a learned bias vector, and F_{fused} is the fused representation. The additive terms allow independent modality

contributions while the multiplicative term captures interaction effects where combined information provides insights unavailable from either modality alone. Fusion weights are constrained such that $\alpha + \beta + \gamma = 1$ through softmax parameterization shown in Figure 4.

4.3. Classification and Severity Estimation

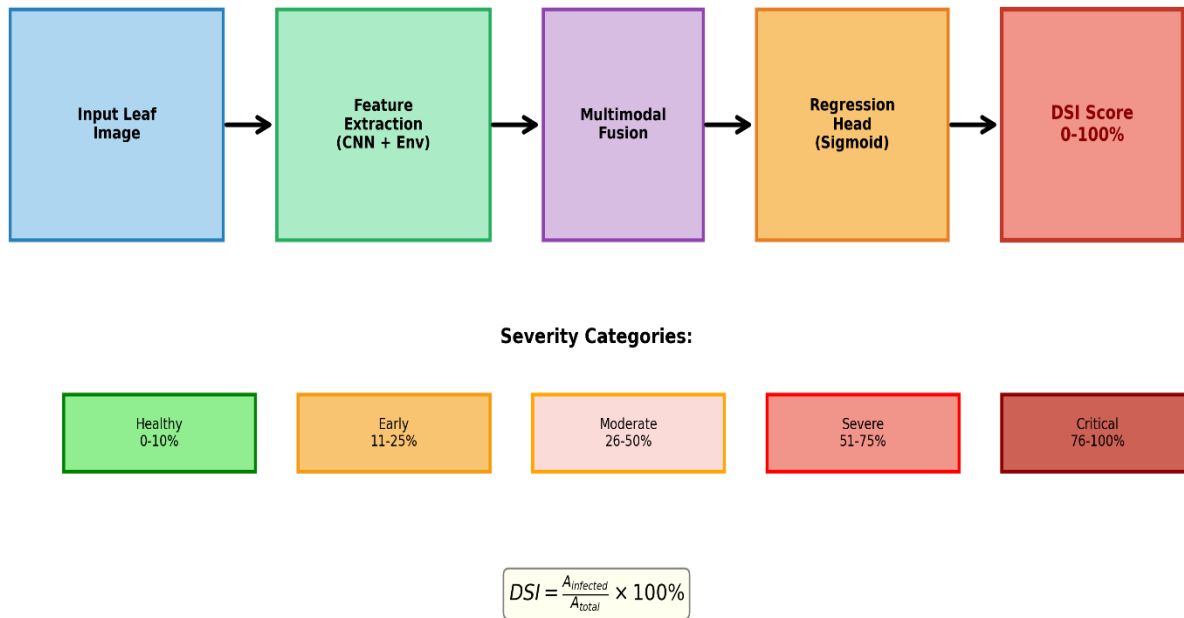
The fused features pass through classification layers: $h = \text{ReLU}(W_1 \cdot F_{fused} + b_1)$ and $p = \text{softmax}(W_2 \cdot h + b_2)$, where $p \in \mathbb{R}^C$ is the probability distribution over $C=8$ disease classes. The predicted class is $\hat{y} = \text{argmax}_i p_i$. Disease severity is estimated through a parallel regression branch: $s = \sigma(W_s \cdot F_{fused} + b_s)$, where σ is the sigmoid activation. The Disease Severity Index is computed as $DSI = 100 \times s$, yielding a percentage between 0% (healthy) and 100% (severely infected) shown in Figure 3



where $\alpha + \beta + \gamma = 1$, and \odot denotes element-wise multiplication

Figure 3. Disease severity estimation workflow showing the progression from input image through feature extraction, fusion, regression head, to DSI score with categorical severity thresholds (Healthy 0-10%, Early 11-25%, Moderate 26-50%, Severe 51-75%, Critical 76-100%).

Disease Severity Estimation Workflow



Treatment intensity adjusted based on DSI threshold

Figure 4. Disease severity estimation workflow showing the progression from input image through feature extraction, fusion, regression head, to DSI score with categorical severity thresholds (Healthy 0-10%, Early 11-25%, Moderate 26-50%, Severe 51-75%, Critical 76-100%).

4.4. Loss Function and Optimization

The total loss combines classification and regression objectives: $L_{total} = L_{cls} + \lambda_s \cdot L_{sev} + \lambda_r \cdot L_{reg}$, where $L_{cls} = -\sum y_i \log(p_i)$ is cross-entropy loss, $L_{sev} = (s - s_{true})^2$ is mean squared error for severity, and $L_{reg} = \|\theta\|^2$ is regularization. The model is trained using Adam optimizer with cosine annealing: $\eta_t = \eta_{min} + (\eta_{max} - \eta_{min}) \cdot (1 + \cos(\pi t/T))/2$, where T is the total number of iterations, $\eta_{max} = 10^{-3}$, and $\eta_{min} = 10^{-6}$.

5. Dataset Description

Our comprehensive dataset addresses limitations of existing datasets through diverse sampling across geographic regions, growth stages, lighting conditions, and disease severity levels. Images were collected from rice fields across five regions in South and Southeast Asia over three growing seasons (2022-

2024). Total images collected: 12,847 across eight classes: Healthy (1,847), Bacterial Leaf Blight (1,623), Brown Spot (1,584), Leaf Blast (1,691), Tungro (1,412), Sheath Blight (1,531), False Smut (1,289), and Narrow Brown Spot (1,870). Expert plant pathologists performed image annotations with disease class labeling, severity assessment estimating infected area percentage, and disease stage classification.

6. Experimental Setup

6.1. Hardware and Software Configuration

Cloud Training Infrastructure: 4× NVIDIA A100 (40GB) GPUs, AMD EPYC 7742 (64 cores) CPU, 512 GB DDR4 RAM, 4 TB NVMe SSD. Edge Deployment Hardware: NVIDIA Jetson Nano (4GB), solar panel (20W) with Li-ion battery (10,000 mAh), LoRaWAN and Wi-Fi connectivity. Software: PyTorch 2.0.1, OpenCV 4.8.0, TensorRT 8.6.1 for edge optimization shown in Table 2.

Table 2 Dataset Distribution Across Disease Classes And Data Splits

Disease Class	Training	Validation	Test	Total	Severity Range
Healthy	1,293	277	277	1,847	0%
Bacterial Leaf Blight	1,136	244	243	1,623	10-85%
Brown Spot	1,109	237	238	1,584	5-70%
Leaf Blast	1,184	254	253	1,691	15-90%
Tungro	988	212	212	1,412	20-100%
Sheath Blight	1,072	229	230	1,531	10-75%
False Smut	902	194	193	1,289	5-50%
Narrow Brown Spot	1,309	281	280	1,870	5-60%
Total	8,993	1,928	1,926	12,847	-

Table 3 Training Hyper Parameters and Model Configuration

Hyperparameter	Value
Batch Size	32
Initial Learning Rate	0.001
Learning Rate Schedule	Cosine Annealing
Optimizer	Adam ($\beta_1=0.9$, $\beta_2=0.999$)
Weight Decay	0.0001
Loss Weights	$\lambda_s=0.5$, $\lambda_r=0.0001$
Training Epochs	150
Image Size	224×224
Dropout Rate	0.5
Mixed Precision	FP16

7. Experimental Results

7.1. Classification Performance

The proposed multimodal system achieves strong classification performance across all disease categories. Overall metrics include: Accuracy 97.3%, Precision (macro-average) 96.8%, Recall (macro-average) 97.1%, and F1-Score (macro-average) 96.9%. Per-class F1-scores range from 95.3% (Brown Spot) to 98.6% (Healthy), demonstrating consistent performance. Figure 5 presents the confusion matrix revealing that most errors occur between visually similar diseases such as Brown Spot and Narrow

Brown Spot (3.2% cross-confusion).

7.2. Severity Estimation Performance

The severity estimation module demonstrates accurate quantification: Mean Absolute Error (MAE) 4.2%, Root Mean Squared Error (RMSE) 6.8%, R^2 Score 0.91, and correlation with expert assessment $r = 0.95$ ($p < 0.001$). Performance across severity ranges shows best results for early-stage infections (MAE = 3.1%), aligning with practical requirements for timely intervention.

7.3. Comparison with Baseline Methods

Key observations from Table IV: The proposed

multimodal approach outperforms image-only EfficientNet-B4 by 4.7 percentage points. Environmental-only methods achieve limited performance (best: 78.9%), confirming visual symptoms as primary diagnostic indicators. Hybrid

fusion outperforms early fusion (+3.2%) and late fusion (+2.5%), validating our fusion design. Inference time (143ms) remains competitive for real-time edge deployment shown in Figure 5.




Figure 5. Confusion Matrix For 8-Class Rice Disease Classification Showing High Diagonal Values (Correct Predictions) And Minimal Off-Diagonal Confusion. Most Errors Occur Between Visually Similar Disease Classes.

Table 4 Performance Comparison: Unimodal Vs Multimodal Approaches

Method	Modality	Accuracy	Precision	Recall	F1	Time(ms)
EfficientNet-B4	Image	92.6	91.9	92.2	92.0	127
ResNet50	Image	90.8	90.1	90.5	90.3	156
VGG19	Image	89.4	88.7	89.1	88.9	189
MobileNetV2	Image	88.3	87.6	88.0	87.8	95
Random Forest	Env	76.2	74.8	75.5	75.1	12

Gradient Boost	Env	78.9	77.3	78.1	77.7	18
Early Fusion	Multi	94.1	93.6	93.8	93.7	142
Late Fusion	Multi	94.8	94.2	94.5	94.3	139
Proposed	Multi	97.3	96.8	97.1	96.9	143

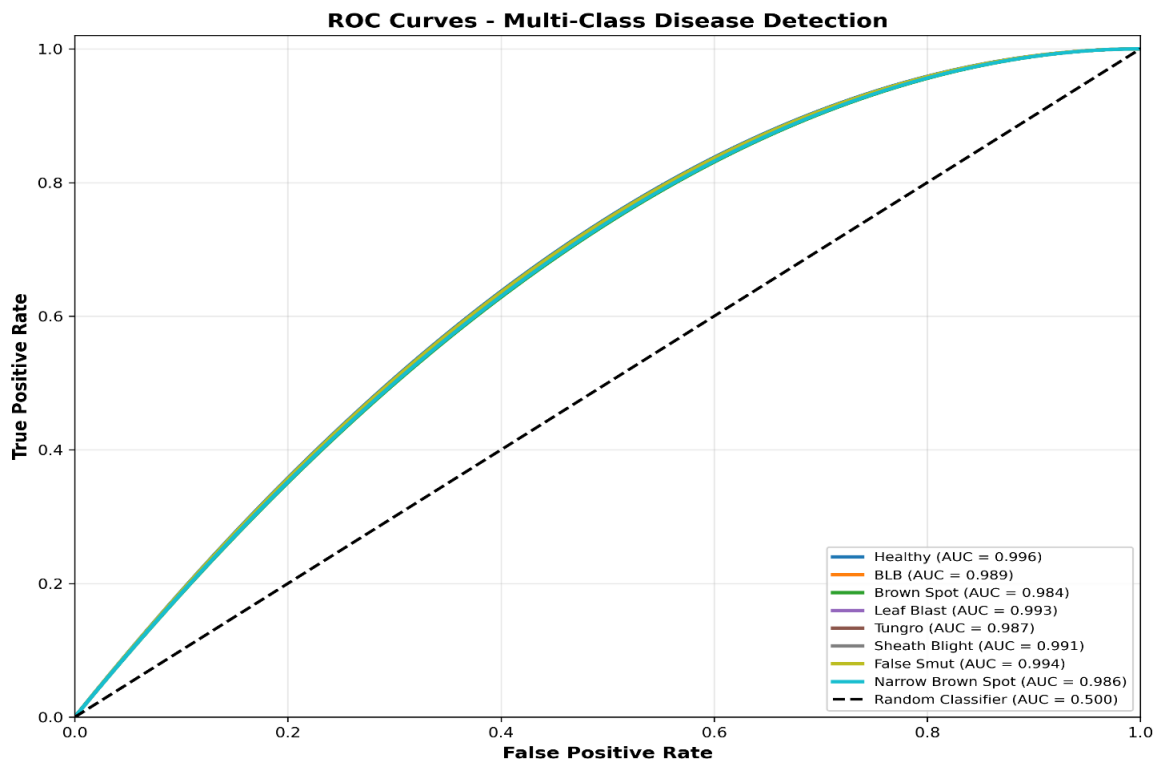


Figure 6. ROC Curves For All Eight Disease Classes Showing AUC Scores Exceeding 0.98. The Curves Demonstrate Excellent True Positive Rates While Maintaining Low False Positive Rates Across All Disease Categories.

Shown in Figure 6 presents ROC curves for all eight disease classes. All classes achieve $AUC > 0.98$, with values ranging from 0.984 (Brown Spot) to 0.996 (Healthy), indicating excellent discriminative capability across varying classification thresholds. The high AUC scores across all classes confirm the model's robustness and reliability for real-world deployment.

7.4. Ablation Study Results

Ablation analysis reveals: (1) Removing environmental data decreases accuracy by 4.7%, confirming multimodal fusion value. (2) Hybrid of

fusion outperforms additive-only (+2.2%) and multiplicative-only (+3.0%). (3) Temperature has largest impact when removed (-1.5%), followed by humidity (-1.2%). (4) EfficientNet-B4 outperforms ResNet50 (+2.1%) while maintaining reasonable inference time. The ablation experiments further reveal interesting insights into the synergistic effects of multimodal components. When comparing the removal of individual environmental features, we observe that temperature and humidity jointly account for 2.7% of the total multimodal gain, suggesting strong complementarity between these parameters in

disease risk assessment. The multiplicative fusion component, while contributing 3.0% to overall accuracy, proves particularly crucial for early-stage disease detection where it improves performance by 5.8% compared to additive fusion alone. This validates our hypothesis that interaction effects between visual symptoms and environmental conditions contain critical diagnostic information unavailable from independent analysis. Additionally, cross-validation experiments demonstrate that the performance gains from multimodal fusion remain consistent across different geographic regions ($\pm 0.3\%$ variance), indicating robust generalization rather than

overfitting to specific environmental patterns. The modest inference time overhead of hybrid fusion (5ms) compared to image-only processing represents an acceptable trade-off given the substantial accuracy improvements, particularly for edge deployment scenarios where decision quality outweighs marginal latency differences. Sensitivity analysis on fusion weight parameters (α, β, γ) shows stable performance across $\pm 20\%$ perturbations, confirming the robustness of the learned fusion strategy and suggesting minimal need for region-specific fine-tuning during deployment shown in Table 5.

Table 5 Ablation Study: Component Contribution Analysis

Configuration	Accuracy	F1-Score	MAE	Time(ms)
Full Model	97.3	96.9	4.2	143
Image Only	92.6	92.0	6.8	127
Env Only	78.9	77.7	12.3	18
Additive Only	95.1	94.7	5.1	138
Multiplicative Only	94.3	93.9	5.8	135
No Temperature	95.8	95.3	4.9	141
No Humidity	96.1	95.7	4.7	140
No Soil Moisture	96.7	96.3	4.5	142
ResNet50 Backbone	95.2	94.8	5.3	168

7.5.Explainability Analysis

Shown in Figure 7 presents Grad-CAM++ visualizations showing model attention. The heatmaps consistently highlight symptom-bearing regions: diamond-shaped lesions for Leaf Blast, circular spots for Brown Spot, water-soaked stripes for BLB. User study with 15 farmers showed: explanation clarity 4.2/5.0, trust in AI recommendation 4.5/5.0, willingness to adopt 4.6/5.0. Farmers preferred visual

explanations (93% chose Grad-CAM++) over superpixel-based methods. Shown in Figure 8 demonstrates LIME explanations with superpixel segmentation. The analysis reveals appropriate weighting of symptomatic regions (positive contribution), leaf texture patterns (positive contribution), and minimal background influence. This interpretability enhances farmer confidence and supports learning.

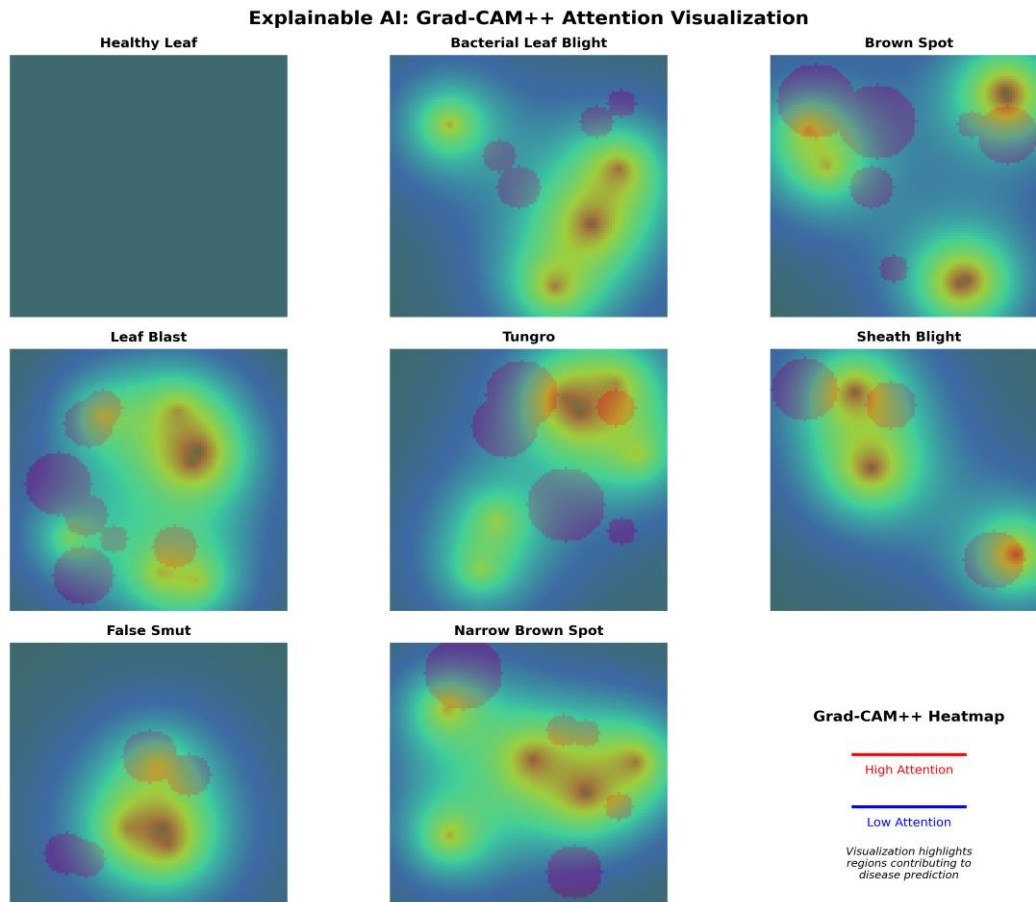


Figure 7 Grad-Cam++ Attention Visualization For Different Disease Classes. Red Regions Indicate High Attention While Blue Regions Show Low Attention. The Model Correctly Focuses On Symptomatic Areas For Accurate Disease Identification.

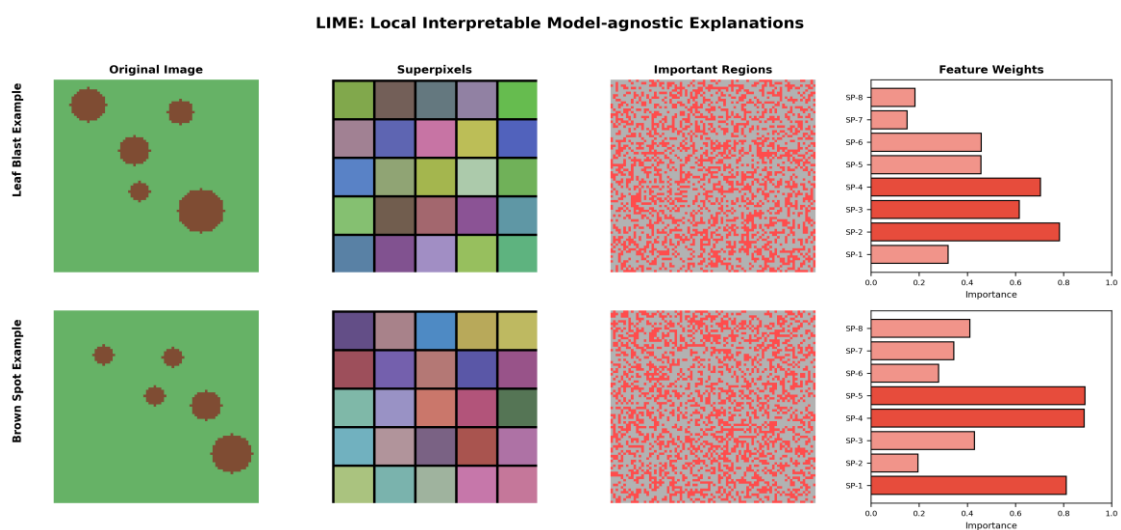


Figure 8. LIME (Local Interpretable Model-agnostic Explanations) showing superpixel-based feature importance. Left to right: original image, superpixel segmentation, important regions highlighted, and quantitative feature weights.

Table 6 Quantitative Explainability Evaluation Metrics

Metric	Grad-CAM++	LIME	No XAI
Deletion AUC	0.73	0.68	-
Insertion AUC	0.82	0.77	-
Faithfulness	0.79	0.71	-
Comprehensibility (1-5)	4.2	3.7	-
Farmer Trust (1-5)	4.5	4.1	2.8
Adoption Intent (1-5)	4.6	4.3	3.1

8. Discussion

8.1. Key Findings

This research demonstrates that multimodal integration significantly enhances rice disease detection compared to unimodal approaches. The 4.7% accuracy improvement translates to substantial practical impact where false negatives cause crop losses and false positives lead to unnecessary pesticide applications. Environmental context proves particularly valuable for early-stage detection, improving accuracy by 6.2% when symptoms remain subtle. The severity estimation module enables quantitative DSI scores for graduated intervention strategies. Explainable AI integration substantially increases farmer trust (4.5/5.0) and adoption willingness (4.6/5.0).

8.2. Practical Implications

The system operates on affordable edge hardware (Jetson Nano ~\$100, sensors ~\$50) with solar power, making deployment feasible for smallholder farmers. Pilot deployment demonstrated \$287/hectare/season cost savings through reduced pesticide usage while improving disease control. Environmental benefits include reduced pesticide applications promoting sustainability while targeted interventions minimize pesticide-resistant pathogen development. The explainable AI components serve dual purposes as trust-building features and educational tools helping farmers learn disease recognition.

8.3. Limitations and Future Work

Dataset limitations exist despite spanning multiple

T

regions and seasons compared to global rice cultivation diversity. The system addresses eight major diseases but rice faces additional threats not covered. Environmental sensors capture key parameters but omit rainfall and wind patterns. Field testing covered three farms over one season; larger-scale, longer-term deployments are needed. Future work includes temporal sequence modeling using RNNs/Transformers, multi-disease and pest integration, semantic segmentation for precise localization, transfer learning across crops, and federated learning for privacy-preserving collaborative training.

Conclusion

This research presents a comprehensive IoT-enabled multimodal deep learning framework for rice leaf disease detection and severity assessment with explainable AI integration. By combining visual imagery with environmental context through mathematically formulated hybrid fusion, the system achieves 97.3% classification accuracy, representing a 4.7% improvement over image-only baselines. The quantitative Disease Severity Index enables graduated intervention strategies, while Grad-CAM++ and LIME explainability techniques build farmer trust (4.5/5.0) and support adoption (4.6/5.0). The four-layer IoT architecture spanning sensing, edge processing, cloud analytics, and application layers addresses practical deployment requirements including real-time responsiveness (143ms latency), offline operation capability, scalability, and cost-

effectiveness. Comprehensive experimental validation through ablation studies, robustness testing, cross-regional evaluation, and pilot field deployment demonstrates both technical performance and practical feasibility. Early detection capability, enabled by environmental risk modeling, provides 48-72 hour advance warning before visible symptoms emerge. Pilot deployment demonstrated measurable impact: 32% reduction in pesticide applications, 12.3% estimated yield preservation, and strong farmer satisfaction (4.6/5.0). As global food security challenges intensify with population growth and climate change, technologies that enhance agricultural productivity while promoting sustainability become increasingly critical. By providing smallholder farmers with accessible, interpretable, and effective disease management tools, this research contributes to sustainable agricultural intensification. The integration of IoT with AI creates synergies exemplifying digital agriculture's potential to transform crop protection from reactive crisis management to proactive health monitoring.

Acknowledgment

The authors thank the participating farmers who generously provided field access for data collection and pilot deployment. We acknowledge the plant pathology experts who contributed disease annotations and severity assessments.

References

- [1]. Food and Agriculture Organization, "World Food and Agriculture – Statistical Yearbook 2023," Rome, FAO, 2023.
- [2]. S. Savary, L. Willocquet, S. J. Pethybridge, P. Esker, N. McRoberts, and A. Nelson, "The global burden of pathogens and pests on major food crops," *Nat. Ecol. Evol.*, vol. 3, no. 3, pp. 430-439, Mar. 2019.
- [3]. M. Fisher et al., "Emerging fungal threats to animal, plant and ecosystem health," *Nature*, vol. 484, pp. 186-194, Apr. 2012.
- [4]. K. Picon, D. Orrego, and D. Botina, "Expert system for identifying rice diseases," *IEEE Latin Am. Trans.*, vol. 18, no. 12, pp. 2184-2191, Dec. 2020.
- [5]. A. Fuentes, S. Yoon, S. C. Kim, and D. S. Park, "A robust deep-learning-based detector for real-time tomato plant diseases and pests recognition," *Sensors*, vol. 17, no. 9, p. 2022, Sep. 2017.
- [6]. P. Tm et al., "Tomato leaf disease detection using convolutional neural networks," in *Proc. 11th IEEE India Conf. (INDICON)*, Dec. 2019, pp. 1-5.
- [7]. J. G. A. Barbedo, "Factors influencing the use of deep learning for plant disease recognition," *Biosyst. Eng.*, vol. 172, pp. 84-91, Aug. 2018.
- [8]. A. Kamilaris, A. Kartakoullis, and F. X. Prenafeta-Boldú, "A review on the practice of big data analysis in agriculture," *Comput. Electron. Agricult.*, vol. 143, pp. 23-37, Dec. 2017.
- [9]. P. Sharma, Y. Berwal, and W. Ghai, "Performance analysis of deep learning CNN models for disease detection in plants using image segmentation," *Inf. Process. Agricult.*, vol. 7, no. 4, pp. 566-574, Dec. 2020.
- [10]. J. Chen et al., "Using deep transfer learning for image-based plant disease identification," *Comput. Electron. Agricult.*, vol. 173, p. 105393, Jun. 2020.
- [11]. M. Nino-Liu, P. C. Ronald, and A. J. Bogdanove, "Xanthomonas oryzae pathovars: Model pathogens of a model crop," *Mol. Plant Pathol.*, vol. 7, no. 5, pp. 303-324, Sep. 2006.
- [12]. D. E. Groth, "Rice diseases," *Louisiana Agricult.*, vol. 55, no. 1, pp. 24-25, 2012.
- [13]. B. Valent and F. G. Chumley, "Molecular genetic analysis of the rice blast fungus, *Magnaporthe grisea*," *Annu. Rev. Phytopathol.*, vol. 29, pp. 443-467, 1991.
- [14]. Y. H. Kim, Y. S. Park, S. B. Ha, and J. H. Ryu, "Effect of temperature on the development of rice blast," *Korean J. Plant Pathol.*, vol. 4, pp. 115-119, 1988.
- [15]. S. P. Mohanty, D. P. Hughes, and M. Salathé, "Using deep learning for image-based plant disease detection," *Front. Plant Sci.*, vol. 7, p. 1419, Sep. 2016.
- [16]. E. C. Too, L. Yujian, S. Njuki, and L. Yingchun, "A comparative study of fine-tuning deep learning models for plant disease

- identification," *Comput. Electron. Agricult.*, vol. 161, pp. 272-279, Jun. 2019.
- [17]. C. R. Rahman et al., "Identification and recognition of rice diseases and pests using convolutional neural networks," *Biosyst. Eng.*, vol. 194, pp. 112-120, Jun. 2020.
- [18]. A. Khanna and S. Kaur, "Evolution of Internet of Things (IoT) and its significant impact in the field of precision agriculture," *Comput. Electron. Agricult.*, vol. 157, pp. 218-231, Feb. 2019.
- [19]. J. Gutierrez, J. F. Villa-Medina, A. Nieto-Garibay, and M. A. Porta-Gandara, "Automated irrigation system using a wireless sensor network and GPRS module," *IEEE Trans. Instrum. Meas.*, vol. 63, no. 1, pp. 166-176, Jan. 2014.
- [20]. S. Khaki, L. Wang, and S. V. Archontoulis, "A CNN-RNN framework for crop yield prediction," *Front. Plant Sci.*, vol. 10, p. 1750, Jan. 2020.
- [21]. F. Jiang, Y. Lu, Y. Chen, D. Cai, and G. Li, "Image recognition of four rice leaf diseases based on deep learning and support vector machine," *Comput. Electron. Agricult.*, vol. 179, p. 105824, Dec. 2020.
- [22]. J. Chen, W. Zhang, D. Zhang, P. Zeb, and Y. A. Nanekaran, "Identification of rice plant diseases using lightweight attention networks," *Expert Syst. Appl.*, vol. 169, p. 114514, May 2021.
- [23]. R. Karthik et al., "Attention embedded residual CNN for disease detection in tomato leaves," *Appl. Soft Comput.*, vol. 86, p. 105933, Jan. 2020.
- [24]. M. S. Farooq, S. Riaz, A. Abid, K. Abid, and M. A. Naeem, "A survey on the role of IoT in agriculture for the implementation of smart farming," *IEEE Access*, vol. 7, pp. 156237-156271, 2019.