An Enhanced Dual-Chain Edge AI Framework for Rice Leaf Disease Diagnosis Under Environmental Degradations

Thanh-Tai Khuu[1, a], Thanh-Nghi Doan[2,b*], Trung-Duc Pham[1,c], Quoc-Anh Nguyen[1,d], Thai-Ngoc Luong[1,e] and Quoc-Vu Nguyen[1,f]

[1]Faculty of Technology and Engineering, Dong Thap University, Vietnam.

[2]An Giang University, VietNam National University, Ho Chi Minh City, An Giang, Vietnam

[a]khuuthanhtai.thcslongkhanhb.hongngu@dongthap.edu.vn, [b]dtnghi@agu.edu.vn, [c]ptduc@dthu.edu.vn, [d]nqanh@dthu.edu.vn, [e]ltngoc@dthu.edu.vn, [f]nqvu@dthu.edu.vn

Keywords: Edge AI, Rice Leaf Disease, Environmental Degradations, Res-UNet & YOLO26m, Static-image Diagnosis.

Abstract. Resource-constrained Edge AI in smart agriculture faces severe performance degradation under environmental and adversarial noise. Addressing this bottleneck, we propose an enhanced dual-chain architecture integrating a Res-UNet defense mechanism with YOLO26m for on-demand static-image diagnosis of four rice leaf diseases. Following a "divide and conquer" strategy, Res-UNet functions as an independent preprocessing shield to restore corrupted images. Thus, YOLO26m detects targets accurately without requiring weight modifications. We empirically validated this approach by expanding 892 rice leaf images into 72 complex noise scenarios. Furthermore, cross-domain benchmarking against 2,000 MS-COCO 2017 and PASCAL VOC 2012 images confirmed the shield's generalization capacity and resistance to overfitting. Quantitative evaluations demonstrate substantial recovery. While severe degradations drop the baseline mean Average Precision (mAP) to 69.7%, our proposed pipeline successfully restores it to 79.8%. Hardware constraint analysis on mobile edge devices reveals a practical inference time of 4018 ms. Ultimately, accepting this latency trade-off to preserve high diagnostic accuracy provides a highly viable strategy for real-world agricultural deployments.

Introduction

Recurring leaf infections—notably Bacterial Blight, Blast, Brown Spot, and Tungro—persistently threaten sustainable rice cultivation, triggering severe economic deficits across developing nations [1-3]. To bypass manual inspections, smart agriculture aggressively shifts toward Edge AI [2, 4]. Deploying lightweight neural networks directly onto resource-constrained embedded IoT sensors eliminates the network dependencies of traditional cloud-based pipelines [4], drastically reducing inference latency to deliver immediate diagnostics [2, 4].

Despite these advantages, mobile edge devices suffer from severe environmental degradations [5] and security vulnerabilities [6]. Uncontrolled optical noise (e.g., sunlight glare, fog, motion blur) ruins image fidelity [5]. Furthermore, attackers inject mathematically precise adversarial perturbations into raw inputs, targeting convolutional extraction layers to deceive the AI into fatal misclassifications [6]. Fusing natural degradation with adversarial noise effectively blinds resource-constrained Edge AI [6].

To solve this bottleneck without forcing the model bloat of conventional adversarial training, we implement a "divide and conquer" dual-chain architecture. A Res-UNet module acts as an independent optical shield, restoring degraded inputs before they reach the YOLO26m detector. This setup neutralizes unpredictable noise without modifying the target classifier's core weights, preserving diagnostic accuracy on lightweight hardware. Rigorous cross-domain benchmarking against MS-COCO 2017 and PASCAL VOC 2012 datasets validates the shield's generalization capabilities and eliminates dataset bias.

Related Work

Agricultural computer vision heavily leverages optimized deep learning architectures—including modified YOLO networks [7], residual attention mechanisms [8], scalable convolutional networks [9], and lightweight inception modules [10]. While highly effective under pristine conditions, these models remain critically exposed to severe optical degradations and adversarial noise. Traditionally, researchers harden systems by injecting mathematically perturbed examples directly into the learning pipeline [11]. Despite providing baseline resilience, this adversarial training introduces severe computational complexities and overfitting risks [12]. Crucially, maintaining noise stability in a monolithic network requires scaling the local Lipschitz constant, bounded by $O(\sqrt{mH})$, which mathematically forces an exponential bloat in parameter count [13]. Deploying such massively bloated adversarial-trained architectures onto ARM-based edge boards triggers inevitable Out-of-Memory hardware failures, violating the absolute memory and power ceilings of agricultural Edge AI [14]. Consequently, providing a direct hardware-level comparison against adversarial training is physically unfeasible under strict edge constraints. To circumvent this limitation, decoupled architectures provide a superior technical pathway by isolating distinct tasks [15]. This isolated processing informs our "divide and conquer" approach. Decoupling the pipeline into an independent restoration shield and a dedicated detection spear actively cleans optical interference without bloating the classifier, securing highly resilient inference speeds on minimal edge hardware.

Proposed Dual-Chain Edge AI Framework

Threat Models and Environmental Degradations: Mobile edge devices in smart agriculture encounter severe optical degradations. Unpredictable environmental interference, including synthetic fog, Gaussian noise, and motion blur, mathematically disrupts the spatial hierarchy of leaf features [5]. To synthesize realistic corruptions, degradation parameters were empirically calibrated against clean ground-truth references: Gaussian noise ($\mu = 0, \sigma = 25$), motion blur (11×11 kernel, $\sigma = 3$), and synthetic fog (35% opacity overlay at intensity 200). Beyond natural stochastic noise, adversarial perturbations present a critical algorithmic vulnerability capable of systematically deceiving object detectors [16]. To evaluate this, we generated adversarial examples using standard Fast Gradient Sign Method (FGSM) [16] and Projected Gradient Descent (PGD) ($\alpha = 0.01$, 10 iterations) [17] protocols. These methods aggressively blind baseline detectors by generating perturbations bounded by intensity $\epsilon = 0.05$. In high-dimensional vector spaces, aligning the noise with the gradient sign ensures every pixel shift pushes neural activations toward the misclassification boundary [16]. This geometric distortion propagates through convolutional layers, drastically shifting bounding box coordinates and blinding the target detector [17].

Res-UNet Defense Mechanism: To neutralize perturbations without altering the target detector, an independent Res-UNet preprocessing shield restores image fidelity [18]. Skip connections prevent the loss of low-level spatial details during downsampling [19], while a pre-trained VGG16 core extracts deep perceptual representations to maintain structural leaf integrity [20]. The complete integration of these decoupled modules is illustrated in Fig. 1.

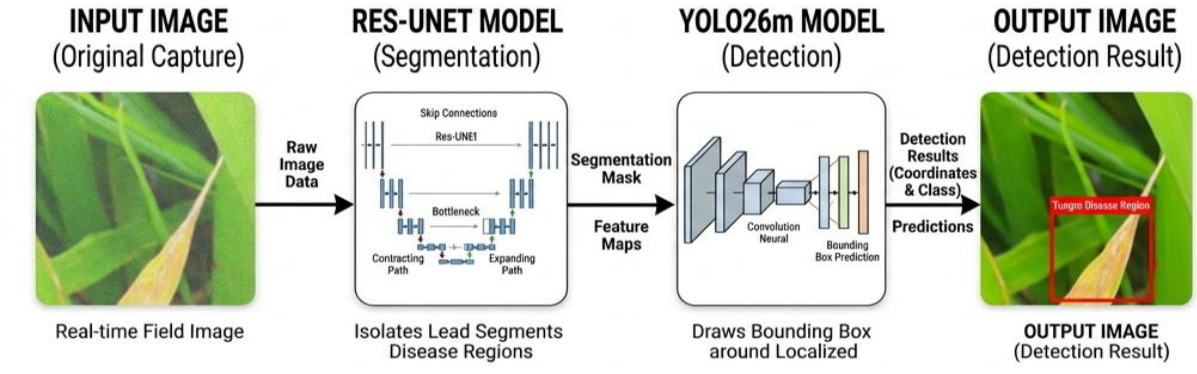


Fig. 1 Architecture Diagram

Balancing pixel-level reconstruction with high-level morphological preservation requires optimizing a generalized objective function [21], as formally expressed in Eq. 1:

$$L_{total} = L_{MSE}(Y, \hat{Y}) + \lambda \cdot L_{perceptual}\left(VGG(Y), VGG(\hat{Y})\right) \quad (1)$$

Here, comparing the clean image Y against the restored output $\hat{Y}$ via VGG feature maps evaluates semantic consistency [18]. We deliberately constrain the perceptual weighting coefficient to $\lambda = 0.1$ [21]. While an overly high coefficient prioritizes abstract style approximation and inadvertently blurs intricate lesion morphologies [19], strictly enforcing $\lambda = 0.1$ compels precise physical boundary reconstruction [18]. This calibration actively prevents superficial optical smoothing [21], directly preserving the crucial high-frequency diagnostic features demanded by the subsequent detection phase [20].

YOLO26m Target Detection: Following structural restoration, the preprocessed input is routed into the YOLO26m architecture, which functions as a highly efficient detection spearhead for the four target leaf diseases. YOLO26 is the latest state-of-the-art vision model officially introduced by Ultralytics [22]. It uniquely eliminates the conventional Distribution Focal Loss (DFL) module in favor of a dual-head architecture for native NMS-free inference and is optimized via the MuSGD algorithm [22, 23]. Crucially, this network remains entirely isolated from the restorative training pipeline. This strict modular encapsulation guarantees that the detector maintains its independent representational integrity, preventing any core weight modifications during noise optimization. While natively engineered for high-throughput execution on resource-constrained devices, we strategically deploy YOLO26m exclusively for on-demand static-image diagnosis, prioritizing absolute diagnostic accuracy over error-prone real-time video processing. To further excel in this environment, the model incorporates Small-Target-Aware Label Assignment (STAL). These architectural enhancements significantly minimize computational overhead and ensure robust feature extraction for tiny lesions, ultimately delivering high-precision anomaly detection for reliable field deployment.

Experimental Setup

Datasets and Test Scenarios: To rigorously benchmark the proposed Res-UNet and YOLO26m pipeline, we construct a multi-domain evaluation matrix comprising 2,892 base images. The primary agricultural target consists of 892 high-resolution rice leaf images categorizing four distinct pathological anomalies. To empirically validate cross-domain generalization and eliminate the risk of task-specific overfitting, we augment the testing phase with 1,000 complex scene images extracted from the MS-COCO 2017 dataset [24], alongside 1,000 diverse object samples from the PASCAL VOC 2012 benchmark [25].

We subject every base image to systematic optical degradation across six stringent attack scenarios. These include targeted algorithmic deception via FGSM and PGD adversarial perturbations, paired with severe environmental corruptions consisting of synthetic fog, motion blur, and Gaussian noise. Additionally, we engineer a Mixed attack scenario that fuses multiple stochastic distortions to replicate unpredictable hardware deployments in harsh field environments. This precise optical degradation strategy generates an expansive testing matrix designed to aggressively challenge baseline detectors. Specifically, each of the 892 base rice leaf images was subjected to 4 environmental corruptions ($892 \times 4 = 3{,}568$) and 8 adversarial scenarios (2 attacks $\times$ 4 target architectures, $892 \times 8 = 7{,}136$). The 2,000 cross-domain base images (MS-COCO and PASCAL VOC) were identically expanded, yielding 12,000 restored instances per dataset. Crucially, all 34,704 generated instances were comprehensively evaluated. The complete dataset distribution and the corresponding validation configurations are explicitly summarized in Table 1.

Table 1. Comprehensive Dataset Distribution and Multi-Domain Evaluation Matrix

Dataset	Base Images	Environmental Corruptions	Target-Specific Adversarial	Restored Instances	Total Physical Images
Rice Leaf Disease	892	3,568	7,136	10,704	22,300
MS-COCO 2017	1,000	4,000	8,000	12,000	25,000
PASCAL VOC 2012	1,000	4,000	8,000	12,000	25,000
Overall Matrix	**2,892**	**11,568**	**23,136**	**34,704**	**72,300**

Implementation Details: To ensure stringent reproducibility, the restorative Res-UNet shield is trained on a high-performance server equipped with dual NVIDIA T4 GPUs using the PyTorch Distributed Data Parallel (DDP) protocol [26]. The framework employs an Adam optimizer with a 1e-4 learning rate, a batch size of 16 (8 per GPU), and concludes training after 50 epochs using a combined loss function (MSE + 0.1 × Perceptual). The training process exhibits a stable convergence trajectory, achieving an optimized loss plateau without inducing optimization oscillations. This stable convergence confirms the DDP framework effectively prevents gradient divergence and synchronization bottlenecks [26], enabling the network to map complex noise distributions consistently. Prior to integration, the YOLO26m target detector was trained on a stratified dataset (70% train, 20% validation, 10% test) for 100 epochs (batch size 16) utilizing the AdamW optimizer (LR 0.01, Cosine Annealing). Spatial and color-space invariances were enhanced via native augmentations, including Mosaic, HSV variations, Affine transformations, and horizontal flipping. Subsequently, the trained defense shield and decoupled YOLO26m detector are deployed onto a Xiaomi K20 Pro edge device (Snapdragon 855 processor) to evaluate real-world inference viability.

Results and Discussion

Impact of Degradations and Visual Artifacts: Qualitative assessments demonstrate the neural network's restorative capacity. As explicitly illustrated in Fig. 2, severe Mixed degradations (fusing synthetic fog, motion blur, and Gaussian noise) obscure chlorotic lesions and destroy spatial hierarchies, natively blinding the detector and erasing bounding boxes. Routing this corrupted input through the Res-UNet shield systematically reconstructs spatial integrity, enabling the YOLO26m architecture to re-establish precise localizations through sharp boundary recovery.



Fig. 2 Diagnostic lifecycle under Mixed corruption

Furthermore, characterizing specific optical disruptions against the pristine baseline, which is illustrated in Fig. 3, reveals that while environmental disturbances degrade global illumination and smooth high-frequency textures, adversarial perturbations scramble pixel neighborhoods without altering macroscopic appearances. Across all disruption profiles, the preprocessing shield actively forces corrupted vectors back into the clean data manifold, preserving the critical morphological features required for accurate bounding box regression.

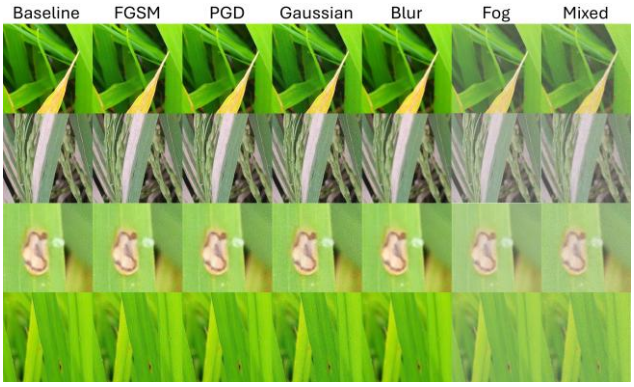


Fig. 3 Comparison of the baseline against distinct adversarial and environmental degradations.

Object Detection Performance: Quantitative assessments establish the structural superiority of YOLO architectures over Torchvision frameworks (Faster R-CNN and RetinaNet) under synthesized degradations. While baseline models suffer severe detection failures-with RetinaNet registering a catastrophic collapse-preprocessing corrupted inputs through the independent Res-UNet shield induces a decisive, cross-architectural performance rebound. As explicitly visualized in Fig. 4, the restored state systematically outperforms the corrupted baseline across key diagnostic metrics. Evaluated against pristine, uncorrupted images, the YOLO26m baseline achieves a clean mAP@0.5 ceiling of 0.9562. Under severe corruption, this metric degrades to 0.697. By effectively intercepting optical perturbations, the Res-UNet shield propels the restored mAP@0.5 back to 0.798. This +0.101 absolute recovery successfully closes approximately 38.96% of the performance gap relative to the clean ceiling. Parallel deterministic recoveries in Precision and Recall confirm the mechanism mathematically reconstructs valid spatial features for bounding box regression rather than merely smoothing superficial noise. Remarkably, this recovery is achieved utilizing a fully frozen YOLO26m detector, confirming that the Res-UNet shield operates beyond a mere denoiser. Governed by the perceptual loss, it functions as a deterministic domain adaptation module that successfully projects corrupted feature vectors back into the clean data manifold. This validates the dual-stage pipeline's operational advantage, which is explicitly quantified in Table 2.

Table 2. YOLO26m performance recovery

Evaluation Metric	Clean Baseline	Corrupted State	Restored State	Absolute Recovery	Gap Closed (%)
mAP@0.5	0.9562	0.697	0.798	+0.101	38.96%
mAP@0.5:0.95	0.8149	0.548	0.646	+0.098	36.72%
F1-Score	0.9121	0.684	0.769	+0.085	37.26%
Precision	0.9017	0.730	0.802	+0.072	41.93%
Recall	0.9226	0.652	0.745	+0.093	34.37%

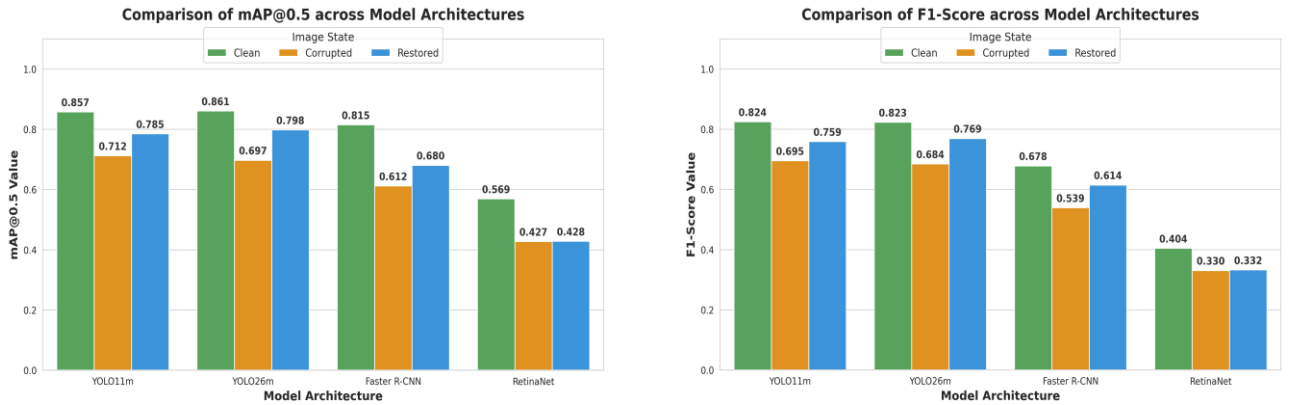


Fig. 4 Cross-architectural comparison of mAP@0.5 and F1-Score

Image Quality Restoration Analysis: We quantify optical recovery using the Structural Similarity Index (SSIM) and Peak Signal-to-Noise Ratio (PSNR), standard benchmarks capturing structural preservation and restoration efficacy [27, 28]. As illustrated in Fig. 5, comparative boxplots expose a distinct transformation. Corrupted rice_leaf populations exhibit severe degradation with extensive interquartile ranges (SSIM standard deviation of 0.218). Post-restoration, distributions demonstrate an upward SSIM median surge to 0.910 coupled with strict variance reduction (standard deviation dropping to 0.026), mathematically proving the network delivers stable enhancements while neutralizing adversarial outliers.

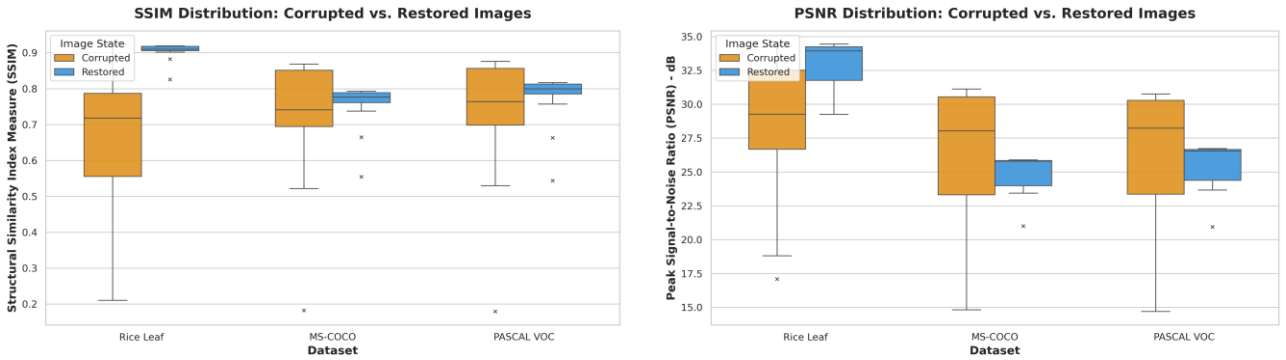


Fig. 5 SSIM and PSNR boxplots comparing Corrupted and Restored image states.

Fig. 6 maps the bivariate correlation between these metrics, segregating the matrix into two non-overlapping clouds. Corrupted instances distribute sparsely within the bottom-left quadrant (minimum SSIM: 0.179), signifying a simultaneous collapse of structural integrity and pixel-level fidelity. Conversely, restored data points soar into a highly concentrated cluster within the top-right quadrant, peaking at 34.44 dB PSNR. This spatial shift confirms operational superiority. Driven by the VGG16 perceptual loss function, the model transcends pixel-smoothing to accurately reconstruct intricate leaf venation patterns and pathological topologies essential for downstream classification.

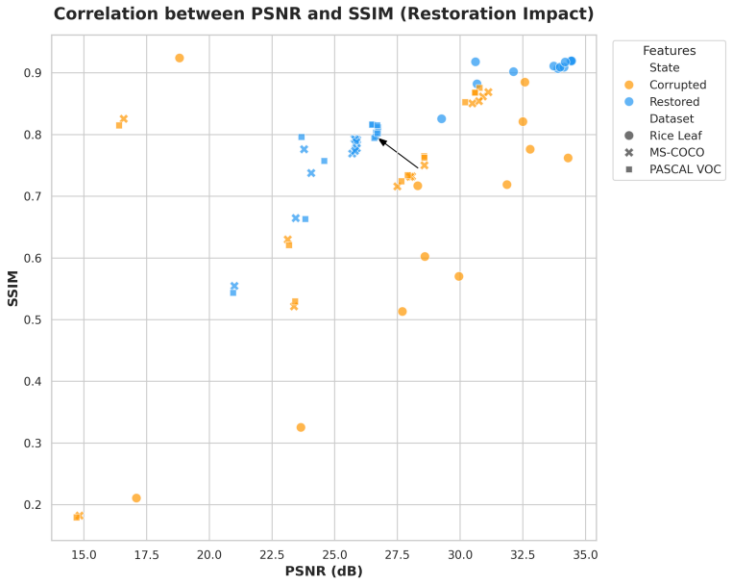


Fig. 6 Bivariate correlation of PSNR and SSIM mapping restoration efficacy.

Hardware Constraints and Deployment Benchmark: To evaluate practical viability, the proposed dual-stage architecture is deployed on a commercial edge device (Snapdragon 855 processor). A deterministic computational trade-off emerges when transitioning from the baseline to the integrated pipeline. The standalone YOLO26m detection head inherently demands substantial

hardware resources, operating at 58.45 GFLOPS with 21.77 M parameters, processing a single frame in approximately 2750 ms. Concatenating the Res-UNet preprocessing shield logically increases the cumulative mathematical operations to 74.7 GFLOPS and 23.71 M parameters. Executing this sequential pipeline on the mobile processor elevates the total inference time to approximately 4018 ms, predictably reducing the overall throughput to 0.25 img/s. These empirical performance metrics, capturing strict hardware limitations, are explicitly summarized in Table 3.

Table 3. Hardware constraints and inference benchmarks on the Xiaomi K20 Pro edge device.

Model	Params (M)	Computational Cost (GFLOPS)	Inference Time (ms)	Throughput (img/s)
YOLO26m (Baseline)	21.7	58.45	≈ 2750	0.37
Res-UNet + YOLO26m	23.71	74.7	≈ 4018	0.25

Despite this elevated latency, the architectural integration guarantees robust detection stability. Fig. 7 explicitly illustrates the practical deployment interface across clean, corrupted, and mathematically restored image states. While the corrupted baseline entirely loses pathological localization, the restored pipeline successfully re-establishes bounding boxes with high diagnostic certainty, registering confidence scores of 0.57 for Brown spot and 0.88 for Tungro. In the context of agricultural pathology, real-time video processing is rarely a mandatory operational requirement. A static diagnostic delay of approximately four seconds per leaf sample falls within the acceptable operational threshold for field agronomists. This constrained temporal increase represents a highly justifiable practical trade-off, sacrificing marginal inference speed to secure peak mAP recovery in harsh, unpredictable environments.

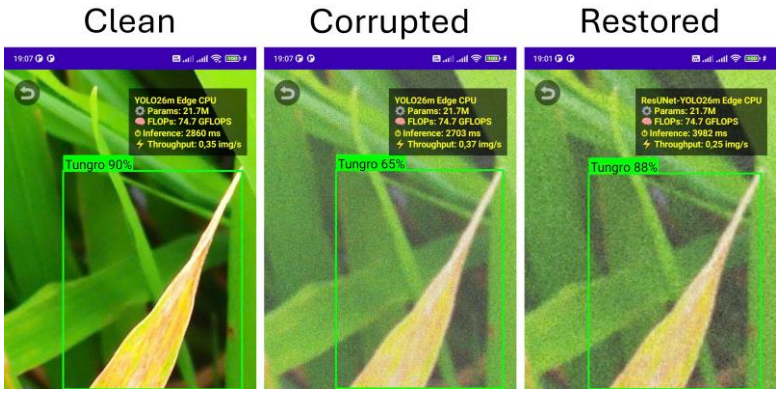


Fig. 7 Real-world mobile deployment interface comparing baseline, corrupted, and restored image states.

Summary

The proposed dual-stage architecture significantly improves the resilience of the diagnostic baseline for rice leaf pathology, maintaining strict cross-domain generalization across heterogeneous external datasets. While achieving peak performance recovery under severe optical degradations, the concatenated pipeline incurs a deterministic computational overhead. To mitigate this latency bottleneck, future research will integrate advanced model quantization and knowledge distillation protocols to aggressively compress the neural network footprint [29]. This approach minimizes the parameter count without compromising the Res-UNet shield's mathematical restoration integrity. Furthermore, transitioning from static devices to autonomous UAV swarm networks represents a critical expansion trajectory. Deploying this edge-optimized framework onto flying platforms will facilitate scalable, continuous aerial surveillance [30], enabling real-time, field-level pathological mapping and decentralized decision-making to secure sustainable precision farming.

References

[1] J. Chen, J. Chen, D. Zhang, Y. Sun, Y.A. Nanehkaran, Using deep transfer learning for image-based plant disease identification, Comput. Electron. Agric. 173 (2020) 105393.

[2] C.G. Simhadri, H.K. Kondaveeti, V.K. Vatsavayi, A. Mitra, P. Ananthachari, Deep learning for rice leaf disease detection: A systematic literature review on emerging trends, methodologies and techniques, Inform. Process. Agric. (2024).

[3] V. Tiwari, R.C. Joshi, M.K. Dutta, Dense convolutional neural networks based multiclass plant disease detection and classification using leaf images, Ecol. Inform. 63 (2021) 101289.

[4] M. Tasci, DBLA-MobileNetV2: Real-Time Rice Leaf Disease Detection on Edge Devices with Dual-Branch Lightweight Attention, IEEE Access 14 (2026) 22356-22371.

[5] C. Fan, J. Zhang, S. Wang, Y. Guo, R. Fu, Robust agricultural pest detection under occlusion and environmental variations via AGIIN-MAF training and SAODL adaptation, Front. Plant Sci. 17 (2026) 1808419.

[6] S.R. Bommana, S. Veeramachaneni, S. Ershad, M.B. Srinivas, Addressing Adversarial Attacks in IoT using Deep Learning AI models, IEEE Access (2025).

[7] D.C. Trinh, A.T. Mac, K.G. Dang, H.T. Nguyen, H.T. Nguyen, T.D. Bui, Alpha-EIOU-YOLOv8: An improved algorithm for rice leaf disease detection, AgriEngineering 6 (2024) 302-317.

[8] J. Peng, Y. Wang, P. Jiang, R. Zhang, H. Chen, RiceDRA-Net: Precise identification of rice leaf diseases with complex backgrounds using a res-attention mechanism, Appl. Sci. 13 (2023) 4928.

[9] M. Tan, Q.V. Le, EfficientNetV2: Smaller Models and Faster Training, arXiv preprint arXiv:2104.00298 (2021).

[10] J. Chen, W. Chen, A. Zeb, S. Yang, D. Zhang, Lightweight inception networks for the recognition and detection of rice plant diseases, IEEE Sensors J. 22 (2022) 14628.

[11] X. Fu, N. Zhou, L. Jiao, H. Li, J. Zhang, The robust deep learning-based schemes for intrusion detection in Internet of Things environments, Ann. Telecommun. 76 (2021) 273-285.

[12] W. Zhao, S. Alwidian, Q.H. Mahmoud, Adversarial training methods for deep learning: A systematic review, Algorithms 15 (2022) 283.

[13] B. Wu, J. Chen, D. Cai, X. He, and Q. Gu, Do Wider Neural Networks Really Help Adversarial Robustness?, arXiv preprint arXiv:2010.01279, (2021).

[14] H. Nugroho, J.X. Chew, S. Eswaran, F.S. Tay, Resource-optimized cnns for real-time rice disease detection with ARM cortex-M microprocessors, Plant Methods 20 (2024) 159.

[15] Z. Zhang, Q. Gao, L. Liu, Y. He, A high-quality rice leaf disease image data augmentation method based on a dual GAN, IEEE Access (2023).

[16] J. Malik, R. Muthalagu, P.M. Pawar, A systematic review of adversarial machine learning attacks, defensive controls, and technologies, IEEE Access 12 (2024) 99421.

[17] A. Thunuguntla, P. Tadepalli, G. Raffa, Defenses against adversarial attacks on object detection: Methods and future directions, Information 16 (2025) 1003.

[18] T. Cevik, N. Cevik, M.S. Ag, O. Guler, S. Saglam, A unified deep learning framework for image denoising and sharpening: from sequential to end-to-end models, Neural Comput. Appl. 38 (2026) 269.

[19] H. Farooq, Z. Zafar, A. Saadat, T.M. Khan, S. Iqbal, I. Razzak, LSSF-Net: Lightweight segmentation with self-awareness, spatial attention, and focal modulation, arXiv preprint arXiv:2409.01572 (2024).

[20] E. Teper, Unified approaches to efficient and robust image retrieval through fusion, restoration, and compression, M.Sc. Thesis, Istanbul Technical University, 2025.

[21] H. Zhao, O. Gallo, I. Frosio, J. Kautz, Loss functions for image restoration with neural networks, IEEE Trans. Comput. Imaging 3 (2017).

[22] G. Jocher, J. Qiu, M. Liu, S. Lyu, F.C. Akyon, M.E. Kalfaoglu, Ultralytics YOLO26: Unified Real-Time End-to-End Vision Models, arXiv preprint arXiv:2606.03748 (2026).

[23] R. Sapkota, R.H. Cheppally, A. Sharda, M. Karkee, YOLO26: Key architectural enhancements and performance benchmarking for real-time object detection, arXiv preprint arXiv:2509.25164 (2026).

[24] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, C.L. Zitnick, Microsoft COCO: Common Objects in Context, in: European Conference on Computer Vision, Springer, Cham, 2014, pp. 740-755.

[25] M. Everingham, L. Van Gool, C.K.I. Williams, J. Winn, A. Zisserman, The PASCAL Visual Object Classes (VOC) Challenge, Int. J. Comput. Vis. 88 (2010) 303-338.

[26] S. Li, Y. Zhao, R. Varma, O. Salpekar, P. Noordhuis, T. Li, A. Paszke, J. Smith, B. Vaughan, P. Damania, S. Chintala, PyTorch Distributed: Experiences on accelerating data parallel training, arXiv preprint arXiv:2006.15704 (2020).

[27] M.G. Azam, A. Kumar, De noise of plant images using generative artificial intelligence and transfer learning, Int. Res. J. Eng. Technol. 12 (2025).

[28] S. Poornima, N. Sripriya, A.F. Alrasheedi, S.S. Askar, M. Abouhawwash, Hybrid convolutional neural network for plant diseases prediction, Intell. Autom. Soft Comput. 36 (2023) 2409.

[29] M.T. Mallick, H. Saha, A. Chakrabarti, Transformative energy-efficient edge-optimised multimodal deep learning framework for pest management and severity analysis in tea plants, J. Agric. Food Res. 15 (2025) 101920.

[30] M.S. Shaikh, B. Kumar, S.K. Gupta, Sustainable smart agriculture using UAV-enabled flying edge computing, J. Syst. Archit. 117 (2025) 102116.