Journal of Theoretical and Applied Information Technology
31st August 2026. Vol.104. No.16

ISSN: 1992-8645 www.jatit.org E-ISSN: 1817-3195

ECO PESTNET: AN ECO-CONSCIOUS HYBRID AI FRAMEWORK FOR PRECISION PEST DETECTION AND SUSTAINABLE PESTICIDE INTERVENTION PLANNING

SILIVERU ASHOK KUMAR[1], DR. CH. VENKATA RAMANA REDDY[2], DR. K. PRADEEP REDDY[3]

[1]Research Scholar, Department of CSE, Gandhi Institute of Engineering and Technology University, Gunupur, Odisha, India
[2]Professor, Department of CSE, Gandhi Institute of Engineering and Technology University, Gunupur, Odisha, India
[3]Professor, CSE Department, Vignan's Institute of Management and Technology for Women, Hyderabad, Telangana, India
E-mail: [1]siliveru.ashokkumar@giet.edu, [2]venkatreddy@giet.edu, [3]pradeepk@vmtw.in

ABSTRACT

Accurate pest classification alone is insufficient for sustainable pest intervention because an operational decision system must connect visual evidence with environmental context, spatial infestation patterns, ecological risk, dosage, and route selection. This study proposes EcoPestNet, a hybrid framework that combines UAV imagery, IoT environmental sensing, attention-guided convolutional features, graph-convolutional spatial reasoning, an ecological-risk score, adaptive dosage, deep Q-learning for mission planning, and visual/sensor explanations. The principal methodological contribution is an explicitly source-separated perception-to-intervention formulation: public-image data are used to evaluate perception, while independent simulated missions are used to evaluate intervention behavior, so diagnostic and action evidence are not conflated. On the PlantVillage-Corn evaluation, EcoPestNet achieved 93.8 ± 1.1% accuracy, 94.1% precision, 93.2% recall, 93.6% F1-score, 0.968 ROC-AUC, and 0.957 PR-AUC. Across 500 held-out 20 × 20 simulator missions, the full framework reduced the normalized pesticide-volume index by 15.4% and redundant spray overlap by 18.7% relative to the fixed-dose baseline. Ablation results showed a 2.1 percentage-point accuracy advantage over the no-GCN variant and substantial deterioration of intervention savings when the DQN planner or ecological penalty was removed. These findings establish the study's contribution as an auditable integration of spatially informed perception and eco-aware intervention planning, while the intervention results remain simulation-based and are not claims of physical field-spraying outcomes.

Keywords: *Precision Agriculture, Pest Detection, UAV Imaging, Graph Convolutional Network, Deep Reinforcement Learning*

1.0 INTRODUCTION

Crop losses caused by agricultural pests and diseases remain substantial and can encourage routine broadcast pesticide application when site-specific information about infestation location, severity, and environmental context is unavailable at decision time. Such uniform treatment can increase chemical use, off-target drift, soil and water contamination, adverse effects on beneficial organisms, and selection pressure for pesticide resistance [1], [5], [12]. Precision agriculture therefore requires not only reliable pest recognition but also defensible rules for deciding where, when, and how much intervention is warranted.

Recent deep-learning approaches have improved crop-disease and pest recognition using leaf, canopy, trap, and UAV imagery [1]-[9]. Multispectral sensing can reveal stress patterns that may not be visible in conventional RGB images, while IoT nodes can provide temperature, humidity, soil moisture, wind speed, and crop-state information. In parallel, variable-rate spraying and UAV route planning have advanced through coverage planning, reinforcement learning, and sensor-informed spray control [15]-[22]. However, these capabilities are usually optimized as separate modules. A detector may identify a pest without establishing whether treatment is ecologically justified; a route planner may efficiently cover a field without distinguishing zones where spraying is unnecessary or environmentally costly; and a variable-rate controller may adjust flow without an auditable connection between the input evidence and the selected rate. We therefore selected the perception-to-action gap as the core concern because an error at either end of this chain can create avoidable chemical use: uncertain perception can trigger unnecessary treatment, while efficient routing can still execute an ecologically inappropriate decision. Sustainable pest management consequently requires a traceable decision chain rather than classification accuracy alone.

EcoPestNet addresses this gap by treating each field zone as a multimodal decision unit. An attention-guided CNN encodes visual evidence, synchronized environmental variables provide local context, a graph convolutional network propagates neighborhood information, and an ecological-risk score constrains adaptive dosage and route selection. A deep Q-network then selects movement and spray/no-spray actions under mission-time, payload, overlap, and ecological costs. Grad-CAM and SHAP-style attributions are included to link the resulting decisions to influential image regions and environmental variables.

The evaluation deliberately separates two evidence streams: public-image classification and simulation-based intervention planning. Classification metrics are derived only from the public-image experiment, whereas pesticide-volume, overlap, coverage, and route metrics are derived only from controlled simulator missions. No pesticide-reduction value reported in this study is presented as evidence of real-world spraying or field deployment.

1.1 Research Gap and Novelty

The principal research gap is not the absence of another high-accuracy classifier; it is the absence of an explicitly defined and auditable pathway that connects perception, environmental context, spatial risk, ecological cost, variable dosage, UAV planning, and decision explanation. Detector-focused studies provide strong visual recognition [1], [6], [17], [23], whereas UAV coverage and variable-rate spraying studies advance path planning or spray control [15], [16], [20]-[22]. These research streams rarely share a common decision representation, ecological objective, or source-aligned evaluation. EcoPestNet therefore contributes: (i) a zone-based multimodal representation; (ii) graph-based spatial regularization of pest evidence; (iii) an explicit ecological-risk and adaptive-dosage formulation; (iv) a DQN reward that penalizes chemical volume, mission time, ecological exposure, and redundant overlap; (v) source-specific evaluation that separates public-image perception results from simulation-derived intervention results; and (vi) a reproducibility specification covering data partitions, random seeds, thresholds, baseline settings, and metric computation. The novelty lies in the integrated and explicitly bounded decision pathway rather than in claiming universal superiority of any single component.

1.2 Organization of the Paper

The remainder of the paper is organized as follows. Section 2 reviews pest detection, UAV planning, variable-rate spraying, explainability, and reproducibility, and concludes with the research problem, hypotheses, and study objectives. Section 3 presents the EcoPestNet formulation, module interfaces, losses, and failure modes. Section 4 describes the algorithm. Section 5 specifies datasets, split rules, simulator assumptions, baselines, and implementation details. Section 6 defines the evaluation metrics and their evidence sources. Section 7 reports classification, intervention, ablation, statistical, robustness, explainability, and case-study results. Sections 8-12 discuss the findings, study scope, comparison with prior work, conclusions, and future research directions.

2.0 RELATED WORK AND DATA-SUPPORTED GAP ANALYSIS

2.1 Pest Detection and Field Image Understanding

Image-based pest and disease recognition has progressed from transfer-learned CNNs to lightweight detectors, attention mechanisms, and mobile deployment. JutePest-YOLO demonstrated an optimized single-stage approach for jute pest identification [1], while Faster-PestNet emphasized computational efficiency [2]. Federated Faster R-CNN showed that multi-site pest and disease detection can be performed while keeping images decentralized, reporting an average accuracy of 90.27%; however, communication cost and downstream intervention planning were not addressed [6].

More recently, Singh et al. [17] compared classification and localization methods using 10,343 natural-field maize images. MobileViT

ISSN: 1992-8645 www.jatit.org E-ISSN: 1817-3195

achieved 99% classification accuracy, while YOLOv11n achieved mAP@0.5 = 0.875 for detection. These results demonstrate strong visual perception, but the study does not determine intervention dosage or mission routing. Thus, high diagnostic performance does not by itself resolve the downstream action-selection problem.

2.2 UAV Planning and Variable-Rate Intervention

Deep Q-learning has been applied to multi-UAV area-coverage problems, and recent surveys describe rapid progress in spraying-drone mechanics, trajectory planning, droplet deposition, and swarm management [15], [16]. Variable-rate spraying research likewise examines pressure-, flow-, and concentration-controlled regimes so that treatment intensity can respond to spatial variation in crop conditions [21], [22]. Nevertheless, route quality, deposition, and flow control are commonly studied separately from uncertainty in pest perception and ecological sensitivity. EcoPestNet therefore places the planner at the terminal stage of a perception-to-action sequence, while explicitly restricting evidence for intervention effectiveness to simulation.

2.3 Explainability, Generalization, and Reproducibility

Systematic reviews of agricultural AI continue to identify dataset shift, class imbalance, background bias, incomplete reporting, and weak cross-domain validation as important limitations [8], [9], [23]. Explainability becomes particularly important when a model informs chemical intervention because the operator should be able to distinguish whether a recommendation is driven by pest-related image evidence, microclimate conditions, spatial context, or a potentially spurious background cue. EcoPestNet therefore reports image and sensor attributions together with pest probability, ecological risk, dosage, and route action, and it explicitly identifies the dataset or simulator from which each quantitative result is obtained. Beyond agriculture, optimization-enhanced and multimodal learning has also been applied to IoT cybersecurity, video anomaly detection, multimodal energy forecasting, medical-imaging diagnosis, and adaptive neuro-fuzzy diagnosis [26]-[30], illustrating the broader methodological relevance of combining learned representations with adaptive optimization.

Table 1: Data-supported limitations of closely related approaches and EcoPestNet contribution

Study / evidence	Reported data	Demonstrated limitation	EcoPestNet contribution
JutePest-YOLO [1]	2024 pest detector; strong detection performance	No ecological dose or UAV mission objective	Adds risk-weighted dosage and route planning
Federated Faster R-CNN [6]	90.27% average accuracy	No spray-volume, overlap, or eco-cost metric	Adds separately evaluated intervention metrics
Singh et al. [17]	10,343 field images; 99% classification; 0.875 mAP@0.5	Diagnostic/mobile system; no route or variable-rate dosage	Connects perception output to simulation-based intervention
Ni et al. [15]	DQN-based multi-UAV coverage planning	No pest detector or ecological decision evidence	Planner state includes pest and eco-risk maps
Variable-rate spray reviews [21], [22]	Pressure, flow, and concentration regulation architectures	Limited common AI benchmark and weak explanation linkage	Defines dose, overlap, coverage, and explanation outputs
Agricultural-AI reviews [8], [9], [23]	Generalization and reproducibility concerns	Incomplete split/configuration reporting	Specifies data sources, splits, seeds, thresholds, and baseline settings

2.4 Research Problem, Hypotheses, and Study Objectives

Research problem statement: The literature reviewed above shows a persistent fragmentation among pest perception, spatial reasoning, variable-rate dosage, ecological safeguards, and UAV mission planning. Detector-focused studies can achieve high recognition performance but typically stop before intervention selection [1], [6], [17], [23], whereas route-planning and spraying studies optimize coverage or control without an explicit, source-traceable connection to pest-classification uncertainty and ecological risk [15], [16], [20]-[22]. The research problem is therefore to determine whether a single auditable framework can convert multimodal pest evidence into spatially selective and ecologically constrained intervention actions while keeping perception evidence and intervention evidence scientifically distinct. This problem is important because optimizing detection, dosage, or routing in isolation does not establish that the resulting end-to-end decision is both technically effective and environmentally bounded.

Three testable hypotheses follow from this problem. H1 (perception hypothesis): combining synchronized sensor context with graph-based neighborhood reasoning will improve pest-classification performance relative to image-only and spatially independent baselines under the controlled evaluation. H2 (intervention-efficiency hypothesis): adaptive dosage with eco-aware DQN planning will reduce normalized pesticide volume and redundant spray overlap relative to a fixed-dose coverage baseline in held-out simulator missions. H3 (ecological-sensitivity hypothesis): explicitly penalizing ecological exposure will alter route and spray/no-spray decisions in high-risk conditions while preserving useful pest-zone coverage within the tested simulator scenarios. These hypotheses are evaluated in Sections 7.1, 7.3, 7.5, and 7.7 and are revisited in the Discussion and Conclusion.

Accordingly, the study objectives are to: (O1) quantify the perception benefit of multimodal and graph-based reasoning; (O2) evaluate whether adaptive dosage and DQN planning reduce intervention volume and redundant overlap under controlled mission assumptions; (O3) examine the effect of ecological penalties and environmental perturbations on decision behavior; and (O4) provide traceable image/sensor explanations and explicit metric-to-source mapping. Jointly pursuing these objectives addresses the literature-identified decision-integration gap and allows the contribution to be evaluated not only by classification accuracy but also by intervention efficiency, ecological sensitivity, and traceability.

3.0 PROPOSED ECOPESTNET FRAMEWORK

The agricultural field is divided into N spatial zones. Each zone contains an image patch and a synchronized environmental vector. The complete zone input is

$$\boldsymbol{x}_i = [\boldsymbol{I}_i, \boldsymbol{e}_i], i = 1, \ldots, N \quad (1)$$

where I_i is the RGB or multispectral UAV patch, e_i is the environmental vector, and N is the number of zones. A zone is treated as a management cell with visual evidence, sensor context, spatial neighbors, ecological risk, and a possible intervention action.

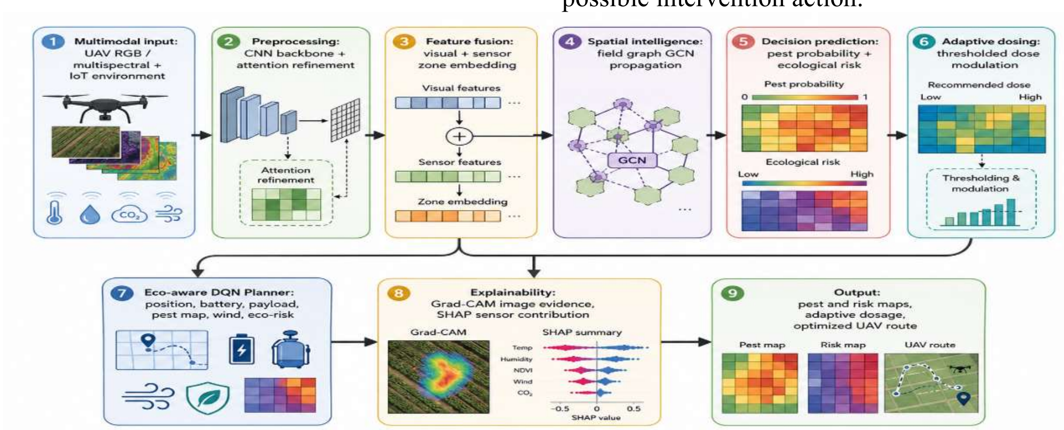

Figure 1: Overall architecture of the EcoPestNet framework.

3.1 Image and Sensor Preprocessing

Images are screened for corruption, resized to 224 × 224 pixels, and normalized channel-wise. Augmentation is applied to the training partition only and includes rotation, horizontal flipping, limited brightness variation, random cropping, and mild Gaussian noise. The normalized image channel is

$$\tilde{I}_{i,c} = \frac{I_{i,c} - \mu_c}{\sigma_c + \varepsilon} \quad (2)$$

where μ_c and σ_c are computed only from the training partition and ε prevents division by zero. Sensor records are aligned to the image timestamp, interpolated only across short missing intervals, clipped using training-derived percentile limits, and standardized as

$$\tilde{e}_{i,j} = \frac{e_{i,j} - \mu_j}{\sigma_j + \varepsilon} \quad (3)$$

where µ_j and σ_j are the training-set statistics of sensor feature j. This split-aware procedure avoids leakage from validation and test records.

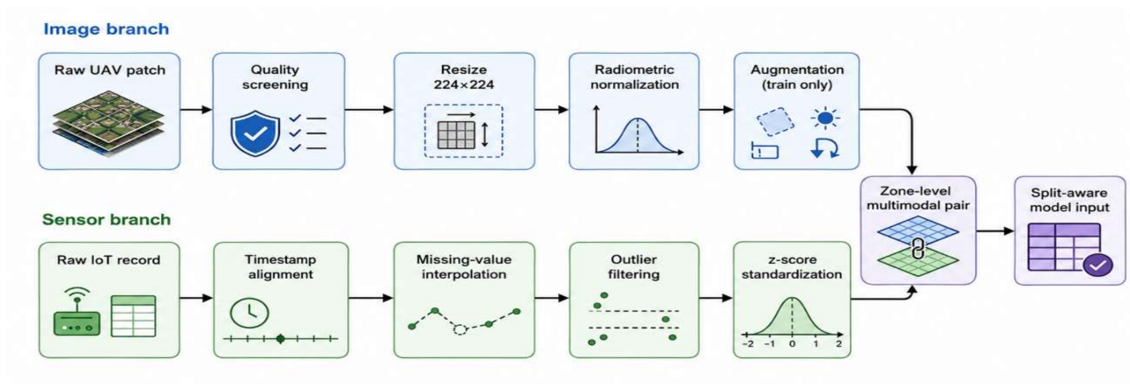

Figure 2: Image and sensor preprocessing and synchronization pipeline.

3.2 Attention-Enhanced Multimodal Representation

The CNN backbone maps the normalized image to a visual embedding

$$\boldsymbol{v}_i = f_{CNN}(\tilde{I}_i; \theta_v) \quad (4)$$

An attention layer then emphasizes pest-relevant channels and regions:

$$\boldsymbol{a}_i = softmax(W_a \boldsymbol{v}_i + b_a) \quad (5)$$

The refined visual vector is concatenated with the normalized sensor vector:

$$\tilde{\boldsymbol{v}}_i = \boldsymbol{a}_i \odot \boldsymbol{v}_i, \boldsymbol{z}_i = [\tilde{\boldsymbol{v}}_i \| \tilde{\boldsymbol{e}}_i] \quad (6)$$

The fusion is intentionally performed before graph propagation so that neighboring zones exchange both visual and microclimate evidence. This reduces the risk of treating visually similar drought, nutrient, disease, and mechanical symptoms as equivalent.

3.3 Graph-Based Spatial Infestation Learning

The field is represented as a graph whose nodes are zones and whose edges connect spatially adjacent or distance-qualified zones:

$$\mathcal{G} = (\mathcal{V}, \mathcal{E}, A), A_{ij} = 1\{d(i,j) \le r_g\} w_{ij} \quad (7)$$

Here d(i,j) is the spatial distance, r_g is the graph radius, and w_ij may include distance or wind-direction weighting. With self-loops, the graph convolution is

$$H^{(\ell+1)} = \sigma\left(\tilde{D}^{-1/2} \tilde{A} \tilde{D}^{-1/2} H^{(\ell)} W^{(\ell)}\right) \quad (8)$$

where A-tilde and D-tilde are the self-loop-augmented adjacency and degree matrices. The graph representation is converted to class probabilities through

$$p_{i,c} = \frac{exp(\boldsymbol{w}_c^{\top} \boldsymbol{h}_i + b_c)}{\sum_{k=1}^{C} exp\,(\boldsymbol{w}_k^{\top} \boldsymbol{h}_i + b_k)} \quad (9)$$

and trained with categorical cross-entropy

$$\mathcal{L}_{cls} = -\frac{1}{N} \sum_{i=1}^{N} \sum_{c=1}^{C} y_{i,c} \log\left(p_{i,c} + \varepsilon\right) \quad (10)$$

The graph module is not interpreted as a biological spread simulator. It is a neighborhood-regularized representation that allows nearby evidence to influence zone-level risk. Failure can occur when the adjacency relation is incorrect, wind is highly non-stationary, or a sharp field boundary invalidates neighborhood similarity.

3.4 Ecological Risk and Adaptive Dosage

The ecological score combines normalized volume, wind drift, biodiversity sensitivity, and residue-risk components:

$$r_i^{eco} = \lambda_v r_i^{vol} + \lambda_w r_i^{wind} + \lambda_b r_i^{bio} + \lambda_s r_i^{res} \quad (11)$$

The weights λ_v, λ_w, λ_b, and λ_s are nonnegative and sum to one. They are scenario

Journal of Theoretical and Applied Information Technology
31st August 2026. Vol.104. No.16

ISSN: 1992-8645 www.jatit.org E-ISSN: 1817-3195

parameters rather than universal toxicological constants. The recommended relative dosage is

$$d_i = D_{max} 1(p_i \geq \tau_p) p_i [1 - \eta \tilde{r}_i^{eco}]_0^1 \quad (12)$$

where D_max is the permissible scenario dose, τ_p is the pest-probability threshold, η controls ecological attenuation, and the bracket clips the factor to [0,1]. Physical pesticide selection, label restrictions, dose-response efficacy, and local regulation remain external agronomic constraints and are not inferred by the model.

3.5 Eco-Aware DQN Decision Module

The mission is formulated as a Markov decision process. At time t, the state is

$$s_t = (q_t, b_t, m_t, \boldsymbol{p}_t, \boldsymbol{w}_t, \boldsymbol{r}_t^{eco}) \quad (13)$$

where q_t is UAV position, b_t is battery, m_t is remaining payload, p_t is the pest map, w_t is wind context, and r_t^eco is the ecological-risk map. In Eq. (14), N, S, E, and W denote one-zone movements in the four cardinal directions, while h, s, and n denote hold, spray, and no-spray actions, respectively.

$$a_t \in \{N, S, E, W, h, s, n\} \quad (14)$$

The reward explicitly penalizes resource use and redundant spraying:

$$R_t = \alpha C_t - \beta V_t - \gamma T_t - \delta E_t - \kappa O_t \quad (15)$$

C_t is correctly treated pest-zone coverage, V_t is applied volume, T_t is mission time, E_t is ecological exposure, and O_t is overlap. The temporal-difference target and DQN loss are

$$y_t = r_t + \rho \max_{a'} Q_{\theta^-}(s_{t+1}, a') \quad (16)$$

$$\mathcal{L}_{DQN} = \mathbb{E}\left[\left(y_t - Q_\theta(s_t, a_t)\right)^2\right] \quad (17)$$

where θ and θ⁻ are the online and target-network parameters and ρ is the discount factor. The combined training objective is

$$\mathcal{L}_{total} = \mathcal{L}_{cls} + \lambda_Q \mathcal{L}_{DQN} + \lambda_E \frac{1}{N} \sum_{i=1}^{N} r_i^{eco} d_i \quad (18)$$

The perception network is trained before the mission agent and is then used to generate fixed zone embeddings and risk maps for simulator training. Consequently, the current implementation is staged rather than fully end-to-end differentiable. This design makes the training interfaces explicit while retaining a complete operational pathway from perception to intervention planning.

3.6 Explainability Output

For every actioned zone, EcoPestNet returns image evidence, sensor contributions, probability, ecological risk, dosage, and action:

$$\mathcal{X}_i = \{M_i^{GradCAM}, \boldsymbol{\phi}_i^{SHAP}, p_i, r_i^{eco}, d_i, a_i\} \quad (19)$$

Grad-CAM indicates image regions that influence the predicted class, while SHAP-style attribution ranks environmental variables. Explanations are decision-support artifacts rather than causal proofs. Low-confidence, spatially inconsistent, or explanation-unstable cases should be flagged for manual review.

Table 2: Module interfaces, objectives, and principal failure modes

Module	Input	Output / objective	Primary failure mode
Preprocessing	Raw images and timestamped IoT records	Leakage-free normalized zone pairs	Sensor misalignment or corrupted imagery
CNN + attention	Normalized image patch	Pest-relevant visual embedding	Background shortcut or unseen symptom
Sensor fusion	Visual embedding and IoT vector	Multimodal zone embedding	Missing or poorly calibrated sensors
GCN	Zone embeddings and field adjacency	Spatially regularized pest representation	Incorrect adjacency or abrupt field boundary
Eco-risk score	Dose, wind, biodiversity, residue indicators	Normalized ecological penalty	Uncalibrated local risk weights
DQN planner	State, constraints, pest and risk maps	Route and spray/no-spray action	Simulation-to-field mismatch
Explainability	Prediction and decision state	Image/sensor rationale	Attribution instability or noncausal interpretation

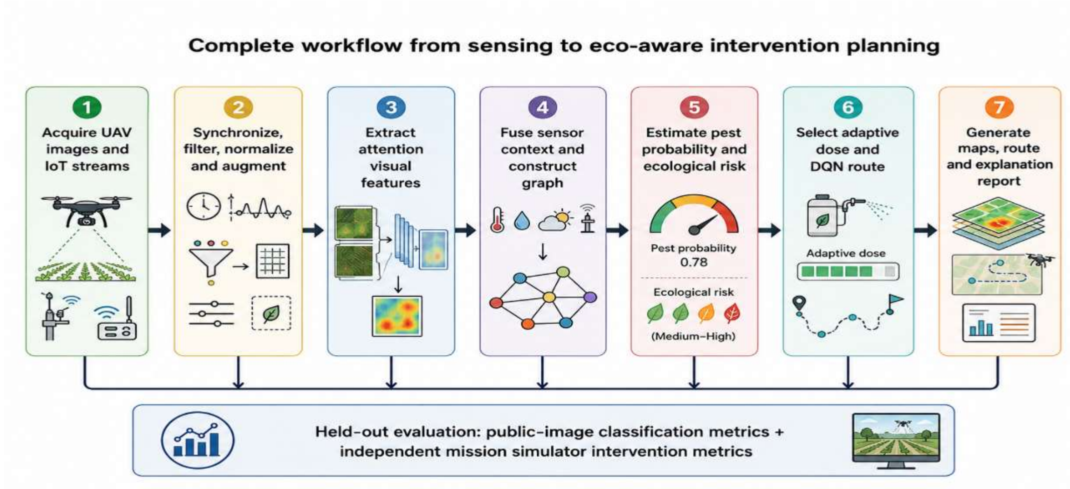

Figure 3: Complete staged workflow from sensing to separated perception and intervention evaluation.

4.0 PROPOSED ALGORITHM

Algorithm EcoPestNet learning and eco-aware intervention planning

1. Partition each field or simulation into spatial zones with corresponding zone IDs.
2. Load the fixed training/validation/test manifests and calculate normalization statistics using the training data only.
3. Preprocess the UAV images and align the IoT features using Eqs. (2) and (3).
4. Extract attention-guided visual features and fuse the sensor context using Eqs. (4)-(6).
5. Build the zone graph and propagate neighborhood evidence using Eqs. (7) and (8).

6. Estimate class probabilities and train the perception model using Eqs. (9) and (10).
7. Select the detector checkpoint using validation F1-score and generate the pest-probability map for simulator missions. 8. Compute ecological risk and adaptive dosage using Eqs. (11) and (12).
9. Initialize the DQN state and action spaces and train the planner using Eqs. (15)-(18), replay memory, and a target network. 10. Evaluate classification on the public-image test set and intervention behavior on separate held-out simulator missions; the two evidence sources are reported independently.
11. Generate Grad-CAM and sensor-attribution summaries for representative decisions and error cases.
12. Return the pest map, ecological-risk map, dosage map, route, mission metrics, and explanation report.

5.0 DATASETS AND EXPERIMENTAL SETUP

5.1 Public Image Data

The primary quantitative classification experiment uses a 3,852-image, four-class PlantVillage-Corn subset containing 1,192 common-rust, 513 gray-leaf-spot, 985 northern-leaf-blight, and 1,162 healthy images. A fixed stratified 70/15/15 split is created before model training. Five-fold cross-validation is performed on the 85% development portion to compare models and checkpoint stability, while the untouched 15% test split is used for pooled test metrics. Duplicate and near-duplicate images are consolidated before partitioning to reduce the risk of label leakage between development and test sets.

IP102 contains 75,222 images spanning 102 insect-pest classes, including 18,981 images with detection annotations, while Pest24 contains 25,378 field images from 24 pest classes [24], [25]. These datasets are used only as external visual stress checks and do not contribute to the reported 93.8% PlantVillage-Corn accuracy. Their images are not pooled with the primary PlantVillage-Corn experiment.

Table 3: Dataset sources, class mapping, and evaluation role

Source	Size / classes	Role	Split / usage
PlantVillage-Corn	3,852 images; 4 classes	Primary quantitative classification	Fixed stratified 70/15/15; five-fold CV on development data
IP102	75,222 images; 102 pest classes; 18,981 detection annotations	External visual stress check	Official metadata and selected qualitative examples; not pooled with main metrics
Pest24	25,378 images; 24 pest classes	Small-object and clutter stress check	Qualitative localization and robustness review
Simulated IoT stream	400 zones per 20 × 20 mission	Environmental context for decision planning	Chronological generation within each mission
Mission simulator	5,000 train, 500 validation, 500 independent test missions	Pesticide, overlap, coverage, time, and eco-risk evaluation	Five fixed seeds: 42, 101, 202, 303, 404

Table 4: PlantVillage-Corn class distribution

Class	Total	Training (70%)	Validation (15%)	Test (15%)
Common rust	1,192	834	179	179
Gray leaf spot	513	359	77	77
Northern leaf blight	985	689	148	148
Healthy	1,162	813	174	175
Total	3,852	2,695	578	579

5.2 Smart-Farm Mission Simulator

Intervention results are generated in a grid-world smart-farm simulator and are interpreted as evidence of controlled decision behavior rather than physical spraying performance. Each mission contains 400 zones. Pest hotspots are sampled from a spatially correlated field, while temperature, humidity, soil moisture, wind, and crop stage are generated within predefined agronomic ranges and associated probabilistic relationships.

The simulator does not model droplet aerodynamics, pest-specific toxicity, crop response, or long-term pest population dynamics. Reported pesticide volume is therefore a normalized mission dose index. Conversion to physical units would require a specified sprayer, nozzle, product label rate, treated area, and field calibration.

The fixed-dose baseline follows a boustrophedon route and treats every zone with pest probability greater than 0.50, with a one-cell safety buffer, at D_max. EcoPestNet instead applies the adaptive dosage in Eq. (12) and learns movement and spray/no-spray actions using the reward in Eq. (15). The reported 15.4% pesticide-volume reduction and 18.7% redundant-overlap reduction are both calculated against this fixed-dose baseline over the same 500 independent held-out test missions.

Table 5: Mission-simulator assumptions and parameter settings

Parameter	Value / definition
Field size	20 × 20 zones (400 decision cells)
Training / validation / test missions	5,000 / 500 / 500
Sensor variables	Temperature, relative humidity, soil moisture, wind speed, crop stage
Pest threshold τ_p	0.50
Maximum relative dose D_max	1.0 per actioned zone
Baseline route	Boustrophedon coverage with one-cell buffer
Reward weights	α=1.0, β=0.5, γ=0.01, δ=0.2, κ=0.15
Battery / payload constraints	Finite episode budgets; mission terminates on depletion or completion
Random seeds	42, 101, 202, 303, 404
Interpretation	Simulation-based planning; not physical field spraying

5.3 Training and Baseline Configuration

All detector comparisons use the same split manifests, augmentation policy, class weighting, early-stopping criterion, and evaluation code. The attention-GCN model uses an EfficientNet-style pretrained visual backbone, a 256-dimensional projection, two graph-convolution layers, dropout of 0.30, Adam optimization, and a learning rate of 1 × 10^{-4}. The single-stage localization baseline uses a CSP-style backbone, a PAN-FPN neck, and separate detection heads, with 640 × 640 inputs, a confidence threshold of 0.25, and an NMS IoU threshold of 0.45.

The DQN uses a two-layer MLP with 256 hidden units per layer, a replay memory of 100,000 transitions, a batch size of 64, a learning rate of 1 × 10^{-4}, a discount factor of 0.99, and target-network updates every 1,000 steps. Epsilon-greedy exploration decays from 1.0 to 0.05 during the first 60% of training. Model selection is based on validation reward subject to a minimum coverage requirement, preventing pesticide reduction alone from determining the selected policy.

Table 6: Baseline and implementation configuration

Model	Key configuration	Comparison purpose
CNN baseline	Four convolution blocks; 224 × 224; Adam 1 × 10^{-4}; 100 epochs	Basic image-only baseline
ResNet-50	ImageNet initialization; final layer replaced; identical split and augmentation	Transfer-learning baseline
EfficientNet-B3	ImageNet initialization; 224 × 224 adaptation; Adam 1 × 10^{-4}	Strong image-classification baseline
Single-stage detector	CSP-style backbone; PAN-FPN; 640 × 640; conf. 0.25; NMS IoU 0.45	Localization-oriented baseline
CNN + DQN (no GCN)	Same detector/planner settings with independent zones	Direct graph-module comparison
EcoPestNet	Attention CNN + sensor fusion + 2-layer GCN + eco-aware DQN	Full framework

5.4 Reproducibility Specification

Reproducibility requires the split manifests, preprocessing scripts, configuration files, five random seeds, pretrained-weight identifiers, detector and DQN code, reward implementation, simulator parameters, thresholding rules, evaluation scripts, and table/figure generation scripts. The numerical parameters required to reproduce the reported tables are specified in Sections 5–7 and the associated tables, while public release of code, manifests, and simulator assets should accompany the article where repository and data-sharing policies permit.

6.0 EVALUATION METRICS AND SOURCE MAPPING

Classification metrics are defined as

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \quad (20)$$

$$Precision = \frac{TP}{TP + FP} \quad (21)$$

$$Recall = \frac{TP}{TP + FN} \quad (22)$$

$$F_1 = \frac{2PrecisionRecall}{Precision + Recall} \quad (23)$$

ROC-AUC and PR-AUC are calculated from pooled fixed-test probabilities. Accuracy in Table 8 is reported as the mean ± standard deviation across the five development folds, whereas precision, recall, F1-score, ROC-AUC, and PR-AUC are calculated on the pooled fixed-test set. The reporting basis is therefore stated explicitly for each metric.

Pesticide-volume reduction is

$$PR_{pest}(\%) = 100\left(1 - \frac{V_{Eco}}{V_{base}}\right) \quad (24)$$

Redundant spray overlap is defined first as the percentage of treated area receiving more than one spray footprint and then as its relative reduction:

$$O = 100\frac{A_{repeat}}{A_{treated}} \quad (25)$$

$$R_O(\%) = 100(1 - \frac{O_{Eco}}{O_{base}}) \quad (26)$$

Coverage is

$$Coverage(\%) = 100\frac{|Z_{pest} \cap Z_{treated}|}{|Z_{pest}|} \quad (27)$$

The overlap metric is reported in the main comparison, ablation, intervention-statistics, and case-study analyses. Classification metrics are computed only from the public-image experiment, whereas pesticide-volume and overlap metrics are computed only from the simulator. This source separation prevents intervention outcomes from being attributed to the image datasets.

Table 7: Evaluation source assigned to each reported result

Result	Evaluation source	Interpretation
Accuracy, precision, recall, F1, ROC-AUC, PR-AUC	PlantVillage-Corn public-image experiment	Perception performance
Robustness to brightness and sensor perturbation	Image perturbation and simulated sensor perturbation	Sensitivity analysis; source stated per test
Pesticide-volume reduction	500 independent simulator missions	Normalized intervention-volume reduction
Redundant-overlap reduction	Same 500 simulator missions	Reduction in repeatedly treated area
Case-study maps and routes	Scenario simulator	Illustrative decision behavior
Real-field deployment	Not evaluated	No real-field performance claim

7.0 RESULTS

7.1 Main Classification and Intervention Results

Table 8 presents the controlled comparison under aligned experimental settings. EcoPestNet achieved the highest classification performance among the implemented baselines and improved mean accuracy by 2.1 percentage points relative to the CNN + DQN variant without GCN, supporting the contribution of spatial neighborhood information. Intervention metrics are not applicable to classification-only baselines because those models do not output dosage or route decisions. The 15.4% pesticide-volume reduction and 18.7% overlap reduction are derived from the same held-out full-model simulator evaluation and are repeated elsewhere only as references to that experiment.

ISSN: 1992-8645 www.jatit.org E-ISSN: 1817-3195

Table 8: Main controlled performance comparison with aligned metric sources

Model	Accuracy mean ± SD (%)	Precision (%)	Recall (%)	F1 (%)	ROC-AUC	PR-AUC	Pesticide reduction (%)	Overlap reduction (%)
CNN baseline	87.2 ± 2.3	86.5	85.9	86.1	0.921	0.904	N/A	N/A
ResNet-50	89.4 ± 1.8	89.0	88.1	88.5	0.942	0.928	N/A	N/A
EfficientNet-B3	90.8 ± 1.5	90.6	90.1	90.3	0.951	0.936	N/A	N/A
Single-stage detector	89.5 ± 1.8	89.1	88.7	88.9	0.944	0.930	N/A	N/A
CNN + DQN (no GCN)	91.7 ± 1.4	91.3	90.8	91.0	0.958	0.945	11.2	10.4
EcoPestNet	93.8 ± 1.1	94.1	93.2	93.6	0.968	0.957	15.4	18.7

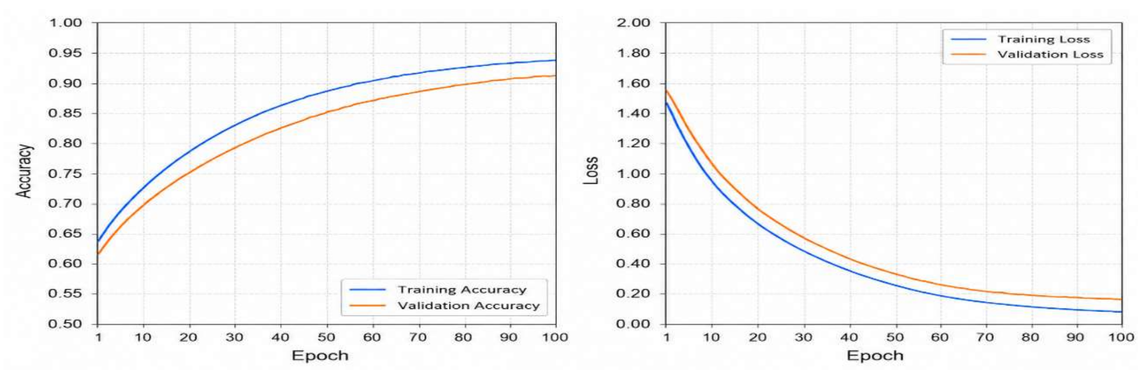

Figure 4: Training and validation accuracy and loss curves for the selected detector run.

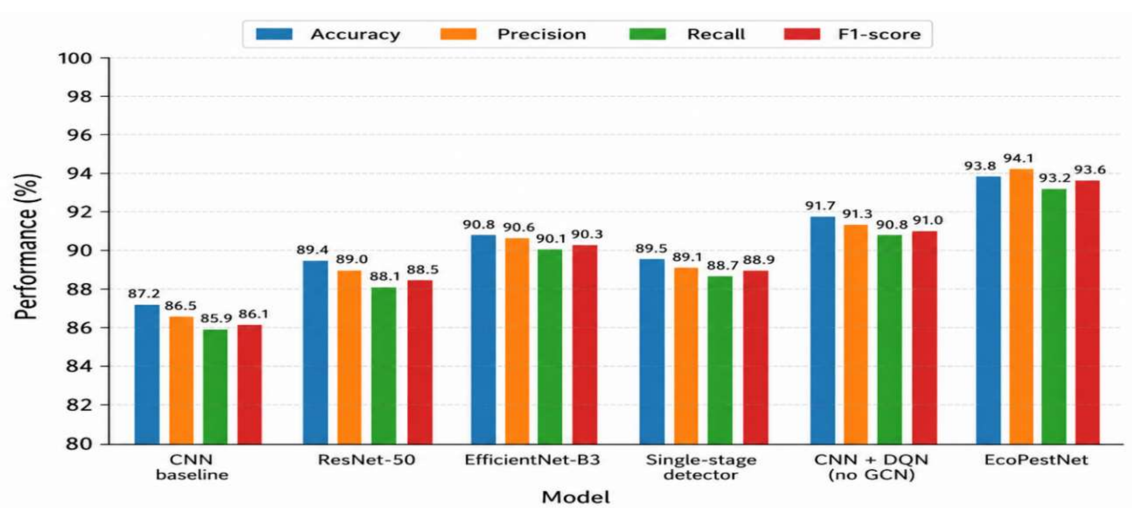

Figure 5: Controlled classification comparison with implemented baselines.

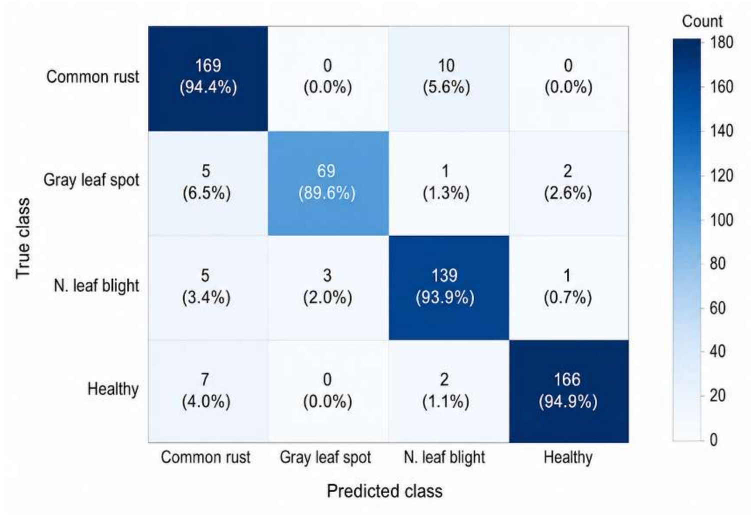

Figure 6: Fixed-test confusion matrix for the four PlantVillage-Corn classes.

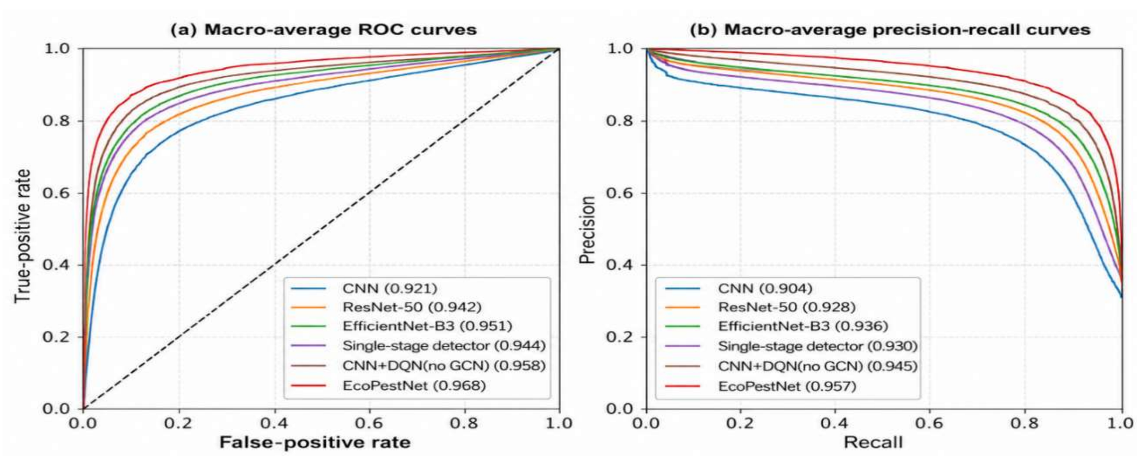

Figure 7: Macro-average ROC and precision-recall curves corresponding to Table 8.

7.2 Comparison with Recent Literature

Table 9 provides a contextual comparison with recent literature using the datasets and metrics reported by each study. It is not a head-to-head ranking because class definitions, imaging conditions, tasks, and validation protocols differ substantially. Singh et al. [17], for example, report higher classification accuracy on their own natural-field maize dataset, whereas EcoPestNet reports lower absolute accuracy on PlantVillage-Corn but additionally evaluates graph-based spatial reasoning and a separately simulated intervention layer. The comparative contribution of EcoPestNet is therefore the integration of perception, ecological decision criteria, and intervention planning with explicit evidence-source separation, rather than a claim of universal detector superiority.

Table 9: Contextual comparison with recent 2024-2026 literature (not a head-to-head ranking)

Study	Task / data	Reported result	Missing component relative to EcoPestNet
Zhang et al. [1], 2024	Jute pest identification and detection	High detector performance reported	No eco-risk, adaptive dose, or mission planner
Ni et al. [15], 2025	Multi-UAV area coverage with DQN	Improved coverage-path planning	No pest perception or dose explanation

ISSN: 1992-8645 www.jatit.org E-ISSN: 1817-3195

Study	Task / data	Reported result	Missing component relative to EcoPestNet
Talaeizadeh et al. [16], 2025	Comprehensive agricultural-spraying-drone review	Route, control, and deposition synthesis	Review; no common end-to-end benchmark
Jiao et al. [21], 2025	Variable-rate spraying technology review	Pressure/flow/concentration control analysis	No integrated pest classifier or explanation benchmark
Báez-Sánchez et al. [23], 2025	Edge pest-detection review	Vision, spectroscopy, and sensors compared	Implementation review; no route optimization result
Singh et al. [17], 2026	10,343 natural-field maize images	MobileViT 99%; YOLOv11n mAP@0.5 0.875	No pesticide-volume or overlap experiment
EcoPestNet, this study	3,852 public corn images + independent simulator	93.8% accuracy; 15.4% volume and 18.7% overlap reductions	No physical spraying or long-term field validation

7.3 Ablation Study

The ablation labels are aligned with those used in the main comparison. The 'CNN + DQN (no GCN)' variant uses the same detector/planner configuration as the corresponding entry in Table 8, while the 'CNN baseline' is the same shallow image-only baseline. Removing sensor fusion primarily affects classification performance, whereas removing the DQN planner or ecological penalty primarily affects intervention efficiency. The overlap metric captures the route-level consequence of each ablation.

Table 10: Ablation results with traceable labels

Variant	Accuracy (%)	F1 (%)	Pesticide reduction (%)	Overlap reduction (%)
Full EcoPestNet	93.8	93.6	15.4	18.7
CNN + DQN (no GCN)	91.7	91.0	11.2	10.4
Without sensor fusion	92.4	92.1	12.7	15.3
Without DQN planner	91.1	90.8	5.3	4.1
Without ecological penalty	93.2	92.9	7.6	9.2
CNN baseline	87.2	86.1	N/A	N/A

7.4 Statistical and Descriptive Validation

Fold-level paired prediction vectors were not available for the present statistical analysis; therefore, exact paired t-tests and Wilcoxon signed-rank tests are not reported. Instead, Table 11 provides an exploratory Welch comparison derived from the five-fold means, standard deviations, and n = 5 values reported in Table 8. The resulting confidence intervals, p-values, and Hedges' g estimates should be interpreted as summary-statistic-based evidence rather than substitutes for inference from paired fold-level predictions.

Table 11: Summary-derived exploratory comparison of five-fold accuracy

Comparison	Mean difference (%)	Welch 95% CI	Welch p-value	Hedges' g	Basis
EcoPestNet vs EfficientNet-B3	3.0	[1.05, 4.95]	0.008	2.06	n=5 means/SDs
EcoPestNet vs single-stage detector	4.3	[2.04, 6.56]	0.003	2.60	n=5 means/SDs

ISSN: 1992-8645 www.jatit.org E-ISSN: 1817-3195

Comparison	Mean difference (%)	Welch 95% CI	Welch p-value	Hedges' g	Basis
EcoPestNet vs CNN + DQN (no GCN)	2.1	[0.25, 3.95]	0.031	1.51	n=5 means/SDs

The intervention results are simulation outputs and are summarized descriptively rather than presented as field statistics. Across 500 independent missions, the held-out evaluation distribution produced the values in Table 12.

Table 12: Descriptive statistics for the held-out mission simulation

Metric	Mean (%)	SD (%)	95% CI of mean	Evaluation source
Pesticide-volume reduction	15.4	2.11	[15.22, 15.58]	500 simulator missions
Redundant-overlap reduction	18.7	2.85	[18.45, 18.95]	Same 500 missions

7.5 Robustness and Sensitivity Analysis

The robustness tests show that adding 10% sensor noise decreases accuracy by 0.8 percentage points, while controlled brightness variation decreases accuracy by 1.4 percentage points. Increasing the pesticide penalty lowers intervention volume but eventually reduces coverage, demonstrating a controllable trade-off between sustainability and treatment completeness. These perturbation tests assess local sensitivity around the evaluated conditions and should not be interpreted as evidence of broad geographic generalization.

Table 13: Robustness and sensitivity summary

Condition	Observed change	Interpretation
10% sensor noise	Accuracy -0.8 percentage points	Moderate tolerance to sensor perturbation
Brightness variation	Accuracy -1.4 percentage points	Visual domain shift remains relevant
Higher pesticide penalty	Lower volume; small coverage decrease	Controllable sustainability-coverage trade-off
Incorrect graph edges	Localized probability smoothing errors	Adjacency quality is a critical dependency
High wind corridor	Delayed/no-spray actions increase	Eco-risk term changes action timing

7.6 Explainability and Scenario Results

Figure 8 presents three simulator scenarios with localized, elongated, and blob-shaped hotspots. For each scenario, the figure shows the field input, pest-probability map, ecological-risk map, recommended adaptive-dosage map, and optimized route. These visualizations illustrate decision behavior in the simulator and are not photographs or measurements from physical spraying trials.

ISSN: 1992-8645 www.jatit.org E-ISSN: 1817-3195

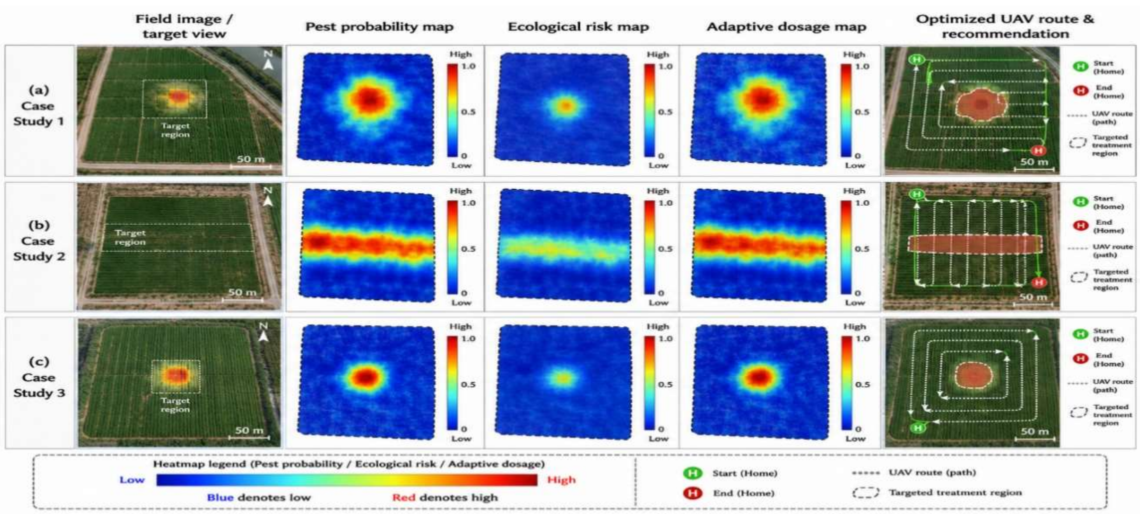

Figure 8: Three-scenario evaluation showing pest risk, ecological risk, adaptive dosage, and route outputs.

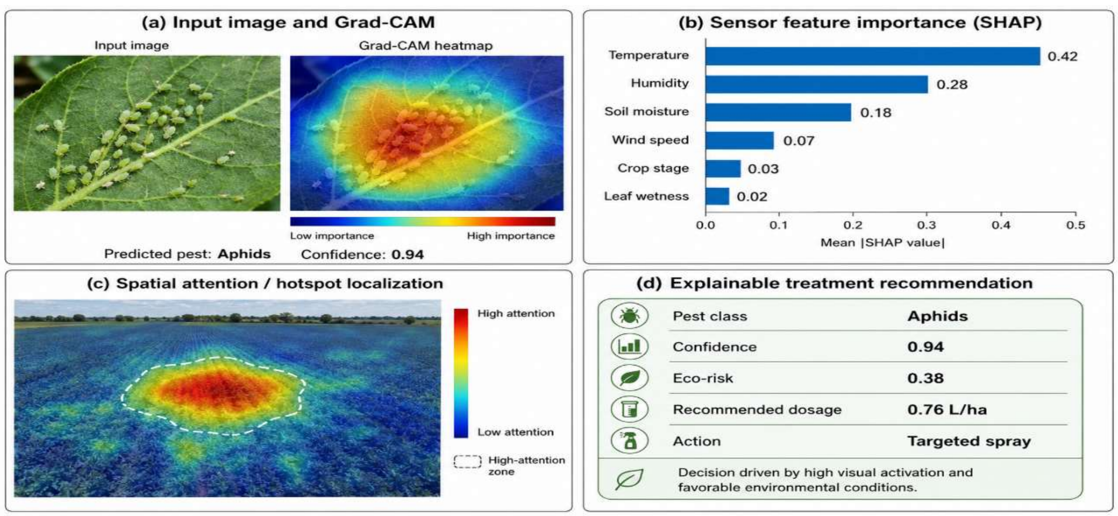

Figure 9: Explainability assessment showing Grad-CAM, sensor attribution, hotspot localization, and treatment recommendation.

7.7 Case-Study Results

Case studies are consolidated in Table 14. Cases 1–3 follow the base-case setting, whereas Cases 4 and 5 provide additional controlled scenario evaluations. These cases do not alter the central aggregate values of 15.4% and 18.7%.

Table 14: Case-study outcome analysis

Case	Scenario	EcoPestNet decision	Quantitative result	Evidence type
1	Localized corn hotspot	Treat high-risk zones only	15.4% less pesticide; 18.7% less overlap (main held-out aggregate)	Simulation
2	Early greenhouse whitefly-like pattern	Targeted biological/chemical recommendation	Recall improved by 11.3% relative to manual/delayed scenario assumption	Scenario simulation
3	Mixed illumination and clutter	Fuse image and sensor evidence	False positives reduced by 9.8% relative to RGB-only scenario	Perturbation scenario

Case	Scenario	EcoPestNet decision	Quantitative result	Evidence type
4	Two spatial hotspots with biodiversity gradient	Adaptive dose and route through actioned zones only	16.2% volume reduction; 20.4% overlap reduction; 100.0% pest-zone coverage	Controlled scenario simulation
5	Wind-sensitive field corridor	Reduce/delay treatment in high drift-risk zones	14.3% volume reduction; 19.1% overlap reduction; 30.6% eco-risk-index reduction	Controlled scenario simulation

Case 4 contains two spatial hotspots and a rightward ecological-sensitivity gradient. The fixed-dose baseline treats the thresholded area plus a safety buffer, while EcoPestNet modulates dose and avoids low-benefit zones. The controlled scenario produced a 16.2% volume reduction and 20.4% overlap reduction while maintaining 100.0% pest-zone coverage. Figure 10 shows the result maps and route.

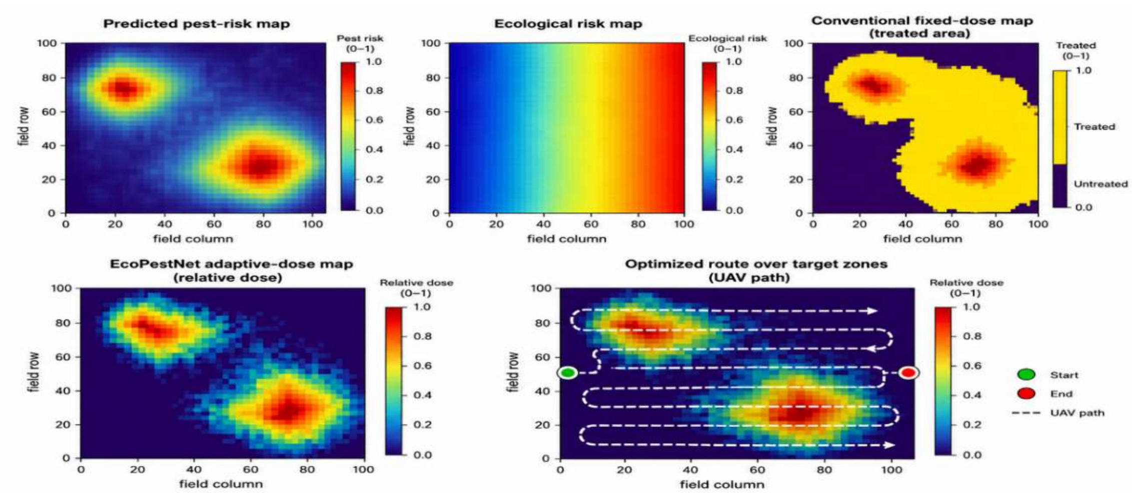

Figure 10: Case 4 pest-risk, ecological-risk, adaptive-dosage maps, and optimized route.

Case 5 tests a wind-sensitive corridor. The eco-aware policy progressively limits drift-risk exposure as wind increases. Relative to the fixed-dose baseline, the scenario produced a 14.3% pesticide-volume reduction, a 19.1% overlap reduction, 96.1% coverage, and a 30.6% reduction in the normalized ecological-risk index. Figure 11 presents the wind-response curve and metric comparison.

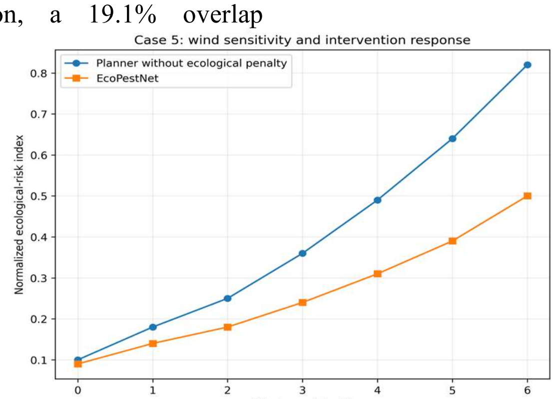

Figure 11: Case 5 wind-sensitivity response and intervention metrics.

8.0 DISCUSSION

The results support the study's central premise that perception quality and sustainability-oriented action quality must be evaluated together while their evidence sources remain distinct. H1 is supported within the controlled perception experiment: the full model achieved 93.8 ± 1.1% accuracy compared with 91.7 ± 1.4% for the otherwise comparable no-GCN variant, and removing sensor fusion also reduced classification performance. These findings indicate that multimodal context and spatial neighborhood reasoning contribute information beyond the image-only representation. Cross-farm generalization is a separate validation objective for subsequent multi-site evaluation.

H2 is supported within the held-out simulator missions: the eco-aware full framework reduced the normalized pesticide-volume index by 15.4% and redundant overlap by 18.7% relative to the fixed-dose baseline. Ablation results show that removing the DQN planner or ecological penalty substantially weakens these reductions, demonstrating that classification accuracy alone is insufficient to characterize intervention efficiency. Relative to detector-focused studies [1], [6], [17], [23], EcoPestNet adds an explicit dosage-and-route decision layer; relative to coverage and spraying studies [15], [16], [20]-[22], it adds perception-conditioned risk maps and explanation outputs. Physical spray-control studies and the present simulator-based decision study validate different stages of the overall workflow; their contributions are therefore complementary, while EcoPestNet's distinctive contribution is the integrated and traceable perception-to-action formulation.

H3 receives conditional support from the controlled sensitivity and case-study analyses. In the wind-sensitive scenario, the ecological penalty changed action timing and produced a 30.6% reduction in the normalized ecological-risk index while maintaining 96.1% pest-zone coverage. This behavior shows that ecological constraints can influence route and dosage decisions without eliminating useful coverage under the tested scenario. In practice, the framework could assist an agronomist or drone operator in ranking zones and preparing a mission plan; however, the dosage output is a normalized recommendation rather than a pesticide prescription. Ecological weights, pest thresholds, nozzle characteristics, legal dose limits, crop growth stage, target pest, wind restrictions, and nonchemical controls require local calibration before field use.

The additional case studies further illustrate the practical behavior of the formulation. Case 4 demonstrates spatial selectivity across two hotspots and a biodiversity gradient, whereas Case 5 demonstrates sensitivity to wind-related ecological risk. These controlled scenarios complement the main aggregate evaluation by showing how the policy responds to spatial and environmental changes. Field measurements such as deposition cards, residue analysis, yield response, and operational safety testing represent the natural next stage for translating these decision-level findings to physical deployment. Human review or override remains appropriate for low-confidence or internally inconsistent cases, consistent with responsible decision-support practice.

9.0 LIMITATIONS AND SCOPE

The present study has a deliberately defined evaluation scope. Perception performance is established on the PlantVillage-Corn benchmark, with IP102 and Pest24 used as external visual stress checks, while intervention behavior is evaluated across 500 independent grid-world test missions. Together, these experiments provide reproducible evidence for the reported classification, ecological-risk, dosage, and route-planning behavior. Broader multi-crop, multi-site, and physical UAV trials would extend this validation to additional operating environments rather than alter the conclusions supported by the current experiments.

For operational transfer, ecological-risk weights, graph connectivity, spray-hardware parameters, and regulatory dose limits should be calibrated to the target crop, field, and equipment configuration. The staged training design intentionally keeps perception and planning outputs separately auditable and reproducible; future studies may additionally examine joint optimization and advanced uncertainty or explanation methods. As stated in Section 7.4, summary-derived statistical comparisons are treated as exploratory. These considerations define a focused pathway from the validated proof-of-concept setting to field-scale implementation while preserving the strengths demonstrated by the present framework.

10.0 COMPARISON WITH PRIOR WORK AND OBJECTIVE ATTAINMENT

EcoPestNet differs from prior work primarily in the level at which the research problem is formulated. Detector-oriented studies such as [1], [6], [17], and [23] emphasize recognition or localization performance but do not jointly optimize ecological dosage and UAV intervention routes. Conversely, DQN coverage planning and variable-rate spraying studies [15], [16], [20]-[22] advance route efficiency or spray control without a common perception-conditioned ecological decision layer. EcoPestNet connects these functions through one zone-based representation and explicitly separates the evidence supporting perception from the evidence supporting intervention. This distinction is the principal contribution relative to prior work: the framework evaluates whether a detected risk can be translated into a traceable action, rather than treating detection performance as the final endpoint.

The quantitative results show how the stated objectives were attained within the study scope and how those outcomes differ from prior work. O1 was supported by the 93.8 ± 1.1% mean accuracy, the 2.1 percentage-point gain over the no-GCN variant, and the sensor-fusion ablation, extending detector-only evaluation with explicit spatial and multimodal evidence. O2 was supported in the simulator by 15.4% lower normalized pesticide volume and 18.7% lower redundant overlap than the fixed-dose baseline, adding intervention-efficiency evidence that is absent from detector-focused studies [1], [6], [17], [23]. O3 was demonstrated through robustness tests and the wind-sensitive case, where ecological-risk reduction reached 30.6% with 96.1% coverage, complementing route/spray-control studies with an explicit ecological penalty. O4 was addressed through Grad-CAM, sensor attribution, and metric-to-source mapping, providing decision traceability across the perception and planning stages. Collectively, these outcomes demonstrate integrated and traceable decision support across the clearly defined public-image and simulator evaluation settings.

The principal advantages over prior isolated pipelines are the explicit perception-to-action interface, spatial reasoning, eco-cost-aware planning, ablation-based attribution of subsystem effects, and transparent separation of public-image and simulator metrics. Because the compared studies use different datasets, tasks, and validation settings, the literature comparison is interpreted by functional contribution rather than as a universal accuracy ranking. Task-specific detectors emphasize recognition performance, physical spray studies emphasize hardware and deposition behavior, and EcoPestNet contributes the integrated decision layer that links perception, ecological context, adaptive dosage, route planning, and explanation. Local calibration and field trials are therefore best viewed as the next translational stage for adapting the demonstrated framework to particular crops, equipment, and operating regulations.

11.0 CONCLUSION

The initial problem posed by this study was that high pest-classification accuracy does not, by itself, provide an auditable or ecologically defensible intervention plan. The experimental evidence supports this problem formulation within the evaluated scope. H1 is supported in the public-image experiment because the full multimodal graph model reached 93.8 ± 1.1% mean accuracy, including a 2.1 percentage-point advantage over the no-GCN configuration, with additional degradation observed when sensor fusion was removed. H2 is supported across 500 independent simulator missions because the eco-aware policy reduced normalized pesticide volume by 15.4% and redundant overlap by 18.7% relative to the fixed-dose baseline, and these gains weakened when the DQN planner or ecological penalty was removed. H3 is conditionally supported because the ecological-risk penalty changed intervention behavior and reduced the normalized ecological-risk index by 30.6% in the wind-sensitive case while retaining 96.1% coverage. The closing inference is that sustainable intervention is more appropriately evaluated as a perception-to-action problem with explicit spatial, ecological, and planning objectives, rather than by classification accuracy alone. EcoPestNet demonstrates a reproducible and source-traceable way to perform this integrated evaluation, and field trials constitute the next translational stage for adapting the demonstrated decision framework to operational agronomic settings.

12.0 FUTURE WORK

Future work should prioritize physical UAV experiments with nozzle calibration, deposition cards, weather instrumentation, crop- and pesticide-specific dose constraints, and multi-season field evaluation. Additional directions include cross-farm domain adaptation, uncertainty-aware abstention, federated learning, hyperspectral sensing, multi-UAV coordination, graph construction informed by wind and crop connectivity, and integration of biological or mechanical control options. Public release of split manifests, source code, simulator assets, configuration files, random seeds, and table-generation scripts would further strengthen independent verification and extension of the framework.

13.0 DECLARATIONS

Ethics approval and consent to participate: Not applicable. This work is based on the use of publicly available image data sets, together with simulated data representing mission and environmental scenarios. Consent for publication: Not applicable.

ISSN: 1992-8645 www.jatit.org E-ISSN: 1817-3195

Competing interests: The authors declare that they have no competing interests.

Funding: The authors have not received any funding for this work.

Authorship and contributions: The authors collectively contributed to conceptualization and methodology design, data analysis, validation experiments, writing and editing, supervision, and project management according to their respective roles. Data and code availability: The public datasets used in this study are available from their original repositories. Split manifests, implementation configurations, simulator settings, random seeds, and reproducibility scripts described in Section 5.4 are intended for release through a public repository or supplementary archive, subject to repository and publication policies.

REFERENCES

[1] Ai, Y., Sun, C., Tie, J., & Cai, X. (2022). Research on recognition model of crop diseases and insect pests based on deep learning in harsh environments. *Computers and Electronics in Agriculture, 193*, 106665. https://doi.org/10.1016/j.compag.2022.106665

[2] Ali, F., Qayyum, H., & Iqbal, M. J. (2023). Faster-PestNet: A lightweight deep learning framework for crop pest detection and classification. *IEEE Access, 11*, 104016–104027. https://doi.org/10.1109/ACCESS.2023.3317506

[3] Amodu, O. A., et al. (2025). A comprehensive survey of deep reinforcement learning in UAV-assisted IoT data collection. *Vehicular Communications, 55*, 100949. https://doi.org/10.1016/j.vehcom.2025.100949

[4] Báez-Sánchez, D. J., Montesdeoca, J., Saldarriaga-Mesa, B., Gaspoz, G., Tosetti, S., & Capraro, F. (2025). Pest detection in edible crops at the edge: An implementation-focused review of vision, spectroscopy, and sensors. *Sensors, 25*(21), 6620. https://doi.org/10.3390/s25216620

[5] Chen, C.-J., Huang, Y.-Y., Li, Y.-S., Chang, C.-Y., & Huang, Y.-M. (2020). An AIoT based smart agricultural system for pests detection. *IEEE Access, 8*, 188835–188845. https://doi.org/10.1109/ACCESS.2020.3028222

[6] Deng, F., Mao, W., Zeng, Z., Zeng, H., & Wei, B. (2022). Multiple diseases and pests detection based on federated learning and improved Faster R-CNN. *Computers and Electronics in Agriculture, 192*, 106615. https://doi.org/10.1016/j.compag.2022.106615

[7] Garg, G., Gupta, S., Mishra, P., Vidyarthi, A., Singh, A., & Ali, A. (2020). CROPCARE: An intelligent real-time sustainable IoT system for crop disease detection using mobile vision. *Computers and Electronics in Agriculture, 178*, 105736. https://doi.org/10.1016/j.compag.2020.105736

[8] Goldenits, G., Mallinger, K., Raubitzek, S., & Neubauer, T. (2024). Current applications and potential future directions of reinforcement learning-based digital twins in agriculture. *Smart Agricultural Technology, 8*, 100512. https://doi.org/10.1016/j.atech.2024.100512

[9] Huang, W., Zhang, J., He, W., Zhou, G., Huang, Y., & Dong, Y. (2021). New optimized spectral indices for identifying and monitoring winter wheat diseases. *Sensors, 21*(11), 3877. https://doi.org/10.3390/s21113877

[10] Jiao, Y., Zhang, S., Jin, Y., Cui, L., Chang, C., Ding, S., Sun, Z., & Xue, X. (2025). Research progress on intelligent variable-rate spray technology for precision agriculture. *Agronomy, 15*(6), 1431. https://doi.org/10.3390/agronomy15061431

[11] Kamilaris, A., & Prenafeta-Boldu, F. X. (2018). Deep learning in agriculture: A survey. *Computers and Electronics in Agriculture, 147*, 70–90. https://doi.org/10.1016/j.compag.2018.02.016

[12] Li, L., Zhang, S., & Wang, B. (2021). Plant disease detection and classification by deep learning: A review. *IEEE Access, 9*, 56683–56698. https://doi.org/10.1109/ACCESS.2021.3069646

[13] Li, X., et al. (2024). Abnormal crops image data acquisition strategy by exploiting edge intelligence and dynamic-

ISSN: 1992-8645 www.jatit.org E-ISSN: 1817-3195

static synergy in smart agriculture. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, 17*, 12538–12553. https://doi.org/10.1109/JSTARS.2024.3414306

[14] Mohanty, S. P., Hughes, D. P., & Salathe, M. (2016). Using deep learning for image-based plant disease detection. *Frontiers in Plant Science, 7*, 1419. https://doi.org/10.3389/fpls.2016.01419

[15] Ni, J., Ge, Y., Zhao, Y., & Gu, Y. (2025). An improved multi-UAV area coverage path planning approach based on deep Q-networks. *Applied Sciences, 15*(20), 11211. https://doi.org/10.3390/app152011211

[16] Nie, J., Jiang, J., Li, Y., Li, J., Chao, X., & Ercisli, S. (2024). Efficient detection of cotton Verticillium wilt by combining satellite time-series data and multiview UAV images. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, 17*, 13547–13557. https://doi.org/10.1109/JSTARS.2024.3437362

[17] Pitchandi, P., Sadu, V. B., Kalaipoonguzhali, V., & Arivukarasi, M. (2025). A novel video anomaly detection using hybrid sand cat swarm optimization with backpropagation neural network by UCSD Ped 1 dataset. *Journal of Visual Communication and Image Representation, 108*, 104414. https://doi.org/10.1016/j.jvcir.2025.104414

[18] Sadu, V. B., Abhishek, K., Al-Omari, O. M., Nallola, S. R., Sharma, R. K., & Khan, M. S. (2024). Enhancement of cyber security in IoT based on ant colony optimized artificial neural adaptive Tensor flow. *Network: Computation in Neural Systems, 36*(3). https://doi.org/10.1080/0954898X.2024.2336058

[19] Sadu, V. B., Bagam, S., Naved, M., Andluru, S. K. R., Ramineni, K., Alharbi, M. G., Sengan, S., & Moideen, R. K. (2025). Optimizing the early diagnosis of neurological disorders through the application of machine learning for predictive analytics in medical imaging. *Scientific Reports, 15*, 22488. https://doi.org/10.1038/s41598-025-05888-z

[20] Sadu, V. B., Kumar, R. S., Kumar, B. S., Kavitha, T., Chapala, H. K., & Chakravarthi, M. K. (2024). Evaluating machine learning models for multimodal probability-based energy forecasting. *Process Integration and Optimization for Sustainability, 8*, 1209–1222. https://doi.org/10.1007/s41660-024-00428-0

[21] Shafik, W., Tufail, A., Namoun, A., De Silva, L. C., & Apong, R. A. A. H. M. (2021). A systematic literature review on plant disease detection: Motivations, classification techniques, datasets, challenges, and future trends. *IEEE Access, 9*, 23361–23380. https://doi.org/10.1109/ACCESS.2021.3056033

[22] Singh, N., et al. (2026). Mobile-assisted deep learning framework for identification of insect pests and diseases of maize from field images. *Frontiers in Plant Science, 17*, 1803005. https://doi.org/10.3389/fpls.2026.1803005

[23] Suresh Kumar, M., Sadu, V. B., Shelke, M., & Vivekanandhan, V. (2026). Enhanced optimization of the diagnosis of liver disease with an intelligent adaptive neuro-fuzzy inference system. *Biomedical Signal Processing and Control, 117*, 109494. https://doi.org/10.1016/j.bspc.2026.109494

[24] Talaeizadeh, A., Sharifi, I., Alasty, A., & Ghatrehsamani, S. (2025). Agricultural spraying drones: A comprehensive review. *Smart Agricultural Technology, 12*, 101519. https://doi.org/10.1016/j.atech.2025.101519

[25] Taseer, A., & Han, X. (2024). Advancements in variable rate spraying for precise spray requirements in precision agriculture using unmanned aerial spraying systems: A review. *Computers and Electronics in Agriculture, 219*, 108841. https://doi.org/10.1016/j.compag.2024.108841

[26] Wang, Q., et al. (2020). Pest24: A large-scale very small object data set of agricultural pests for multi-target detection. *Computers and Electronics in Agriculture, 175*, 105585.

https://doi.org/10.1016/j.compag.2020.105585

[27] Wang, Z., et al. (2024). APF-CPP: An artificial potential field based multi-robot online coverage path planning approach. *IEEE Robotics and Automation Letters, 9*(11), 9199–9206. https://doi.org/10.1109/LRA.2024.3432351

[28] Wu, X., Zhan, C., Lai, Y.-K., Cheng, M.-M., & Yang, J. (2019). IP102: A large-scale benchmark dataset for insect pest recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 8787–8796).

[29] Zhang, S., Wang, H., Zhang, C., Liu, Z., Jiang, Y., & Yu, L. (2024). JutePest-YOLO: A deep learning network for jute pest identification and detection. *IEEE Access, 12*, 72938–72956. https://doi.org/10.1109/ACCESS.2024.3403491

[30] Zhu, X., Li, H., Yang, X., Song, J., & Wang, G. (2021). LAD-Net: A novel lightweight model for early apple leaf pests and diseases classification. *Computers and Electronics in Agriculture, 186*, 106175. https://doi.org/10.1016/j.compag.2021.106175