

AgriSpot-Trans: Fine-Grained Micro-Lesion Segmentation in Complex Unstructured Agricultural Environments via Occlusion-Aware Hierarchical Transformer

Xianglin Qu¹, Haibo Peng^{1,*}, Xuanyi Zhu¹, Xiankang Shen¹, Xinran Chen¹, Jingyun Luo¹ and Rong Zhou²

¹Yunnan Open University, Kunming 650500, Yunnan, China

²Baosight Software(Yunnan) Co., Ltd., Kunming 650500, Yunnan, China

Accurate pixel-level segmentation of early-stage micro-lesions in unstructured agricultural environments is critical for precision agriculture, yet it remains fundamentally challenged by severe field occlusions, optical interferences, and fuzzy biological boundaries. To address these issues, we propose AgriSpot-Trans, a novel end-to-end occlusion-aware hierarchical Transformer. Unlike conventional hard-cropping pipelines that disrupt global canopy topology, our architecture introduces an Environment-Aware Token Routing (E-Decoupler) module. Driven by a Gumbel-Softmax mechanism, it softly suppresses background noise and specular reflections in the feature space. Furthermore, to overcome the morphological mismatch of standard square-window attention, we design Morphological Radial Attention (MRA), which dynamically aligns its star-shaped receptive field with the radial diffusion gradients characteristic of biological tissue damage. Finally, a Level-Set Signed Distance Function (SDF) Decoder, optimized by a Boundary-Aware Active Contour loss, is employed to achieve continuous, sub-pixel contour evolution. Extensive experiments on PlantVillage, RoCoLe, and our newly curated Complex-Field-Pest (CFP) benchmark demonstrate that AgriSpot-Trans achieves a state-of-the-art mIoU of 78.2% in complex field scenarios—a 9.5% absolute improvement over the strong SegFormer baseline ($p < 0.01$). Most notably, it boosts the recall for extreme micro-targets ($= 5$ px) by an unprecedented 23.7%, delivering a highly robust and edge-efficient visual perception framework for real-time, variable-rate pesticide application.

Keywords: Precision agriculture, phytopathology, micro-lesion segmentation, vision transformer, level-set method, unstructured environments, active contour.

1. INTRODUCTION

The early and accurate perception of plant diseases and pests captured by agricultural visual sensors is a cornerstone of modern precision agriculture. Pixel-level annotations of

*Address for correspondence: Haibo Peng, Yunnan Open University, Kunming 650500, Yunnan, China, Email: ^a85005702@qq.com, ^bpenghaibo@ynou.edu.cn, ^c1336832774@qq.com, ^d208869652@qq.com, ^e1603495846@qq.com, ^f2431594080@qq.com, ^g8573746@163.com

micro-lesions, early fungal spots, and pest eggs not only provide indispensable information for automated variable-rate pesticide application (VRA) systems, but also enable the generation of high-fidelity synthetic data to strengthen the robustness of agricultural robots. By treating diseases at their microscopic inception, farmers can drastically reduce chemical usage, ensuring both food security and ecological sustainability.

Nevertheless, executing this task in real-world, unstructured agricultural environments presents formidable challenges. First, unlike well-defined targets in autonomous driving, early-stage phytopathological targets typically appear in images as tiny (often ≈ 5 pixels), morphologically diverse, and sparsely distributed spots. These micro-lesions have faint color gradients that blend seamlessly with healthy chlorophyll, making them easily confused with background noise, sensor dust, or natural leaf pigment variations. Second, in field-based onboard scenarios (e.g., UAVs or ground rovers), cameras capture canopies characterized by extreme unstructured complexity. Targets are unavoidably occluded by overlapping leaves and crossing branches. Furthermore, dynamic weather conditions introduce severe optical interferences, including dappled shadows, strong sunlight washouts, and specular reflections from morning dew[7,8], exponentially increasing the difficulty of accurate annotation.

Traditional, fully convolutional network (FCN)-based semantic segmentation approaches, despite their remarkable success in general scenes and laboratory-controlled agricultural imagery, often underperform on such fine-grained field tasks. Their restricted and fixed receptive fields hinder the capture of global contextual leaf structures—cues that are essential for distinguishing a genuine disease spot from an abiotic stress mark or background artifact. Recently, Vision Transformers (ViTs) and their variants have revolutionized computer vision by capturing long-range dependencies. However, early ViT architectures yield only a single-scale, low-resolution feature map, and their global self-attention incurs quadratic complexity with respect to image size. This renders them computationally useless for pixel-level dense prediction tasks requiring high-resolution inputs, which is mandatory for micro-lesion detection.

To overcome these limitations, hierarchical Transformer architectures—such as SegFormer[1] and Swin Transformer[2]—have emerged. SegFormer attains both efficiency and strong segmentation accuracy via a hierarchical Transformer encoder coupled with a lightweight MLP decoder. Meanwhile, the Swin Transformer introduces the innovative “shifted-window” mechanism, enabling cross-window information exchange while maintaining linear computational complexity.

Inspired by these advances, while being acutely aware of the unique morphological and environmental challenges of phytopathology, we propose a novel two-stage framework specifically tailored to pixel-level disease and pest annotation in complex fields. Our core approach is “localize first, refine second”. In practical terms, we first employ an Illumination & Foliage Decoupling module—a lightweight, occlusion-aware network—to rapidly filter out complex background noise, overlapping physical occlusions, and severe optical interferences, thereby generating clean Regions of Interest

(RoIs) representing the unobstructed leaf surface. Subsequently, within these RoIs, we deploy our more powerful AgriSpot-Trans module. Its encoder is constructed on top of the Swin Transformer but is heavily augmented with morphology-aware attention mechanisms to fully exploit its strengths in capturing the fuzzy, irregular diffusion patterns typical of biological lesions. Furthermore, to address the extreme class imbalance and the ambiguous margins of spreading diseases, we introduce a Boundary-Aware Active Contour loss during the decoding phase, ultimately achieving precise, sub-pixel segmentation of pest and disease spots.

The principal contributions of this paper are:

1. **An Environment-Aware Cascade Framework:** We propose a two-stage “localize first, refine second” architecture that explicitly decouples physical and optical occlusions in unstructured agricultural environments, effectively balancing computational efficiency and micro-target annotation accuracy.
2. **The AgriSpot-Trans Module:** We design a specialized hierarchical Transformer that incorporates morphology-aware self-attention, specifically optimized for the irregular gradients and fuzzy boundaries of early-stage plant diseases.
3. **Boundary-Aware Loss Formulation:** We introduce a novel composite loss function integrating Active Contour principles, which explicitly addresses the extreme foreground-background class imbalance and forces the network to learn the progressive diffusion boundaries of micro-lesions.
4. **Comprehensive Empirical Validation:** We demonstrate the state-of-the-art performance of our method not only on public benchmarks, but also on a newly-curated Complex-Field-Pest dataset, proving its robustness for real-world deployment on edge-computing agricultural devices.

2. RELATED WORK

The accurate pixel-level annotation of micro-lesions in unstructured agricultural fields intersects with several active research domains. In this section, we review the development of deep learning in phytopathology, semantic segmentation in complex agricultural environments, and the recent evolution of Vision Transformers for fine-grained perception.

2.1 Deep Learning in Phytopathology

In recent years, deep learning has fundamentally transformed agricultural disease detection. Early efforts focused mainly on image-level classification using Convolutional Neural Networks (CNNs) such as ResNet, VGG, and MobileNet, demonstrating remarkable accuracy on public datasets like PlantVillage[24,26]. Subsequent research advanced to object detection paradigms (e.g., Faster R-CNN and YOLO series)

to localize infected areas with bounding boxes. However, these approaches have two critical limitations when applied to precision agriculture. First, bounding boxes inevitably capture a substantial amount of healthy tissue and background noise, failing to provide the precise pixel area required for calculating accurate disease severity and guiding Variable Rate Application (VRA) systems. Second, the majority of existing models are trained on lab-controlled imagery with homogeneous backgrounds and detached leaves. When deployed in real-world, unstructured fields, their performance deteriorates sharply due to the domain gap caused by variable illumination, camera motion blur, and complex soil/weed backgrounds.

2.2 Semantic Segmentation in Unstructured Agricultural Environments

To achieve precise disease quantification, researchers have increasingly adopted semantic segmentation architectures. Since the advent of the Fully Convolutional Network (FCN), models such as U-Net and DeepLabv3+ have been widely applied to agricultural tasks such as weed-crop discrimination and canopy extraction. To capture targets at various scales, these networks utilize techniques such as dilated convolutions and Spatial Pyramid Pooling (ASPP). Despite their utility, CNN-based approaches are inherently limited by their localized receptive fields. In complex unstructured environments, early-stage micro-lesions (e.g., initial rust spores or early blight) often occupy fewer than 5 pixels and present extremely weak color gradients against healthy chlorophyll. Furthermore, field images are heavily corrupted by overlapping foliage (physical occlusion) and specular reflections from waxy leaves or morning dew (optical interference). Traditional CNNs struggle to aggregate the global contextual information necessary to distinguish a genuine micro-lesion from a similarly colored abiotic stress mark or a dappled shadow, frequently leading to high false-positive rates.

2.3 Vision Transformers for Fine-Grained Perception

Recently, Vision Transformers (ViTs) have emerged as a powerful alternative to CNNs, leveraging self-attention mechanisms to capture long-range global dependencies. However, standard ViTs employ fixed-size patches and global self-attention, incurring quadratic computational complexity that hampers the processing of high-resolution inputs necessary for dense, micro-target prediction. To address this, hierarchical vision Transformers have been proposed. The Swin Transformer introduces a shifted-window (SW-MSA) mechanism, enabling efficient cross-window feature exchange with linear complexity, thereby producing robust multi-scale representations. Concurrently, SegFormer simplifies the architecture with a positional-encoding-free encoder and a lightweight MLP decoder, achieving an excellent balance of speed and accuracy.

While these advanced hierarchical ViTs have shown promise in general scenarios, their direct application to agricultural micro-lesions remains sub-optimal. Standard Transformers lack the specific inductive biases required to model the fuzzy, irregular diffusion boundaries characteristic of biological tissue damage. Furthermore, the processing of entire images containing severe field occlusions wastes computational resources and degrades feature purity. Our proposed AgriSpot-Trans addresses this shortcoming. By building upon the advanced architectures of SegFormer and Swin Transformer, we introduce a domain-specific “localize first, refine second” cascade framework. We uniquely integrate an Illumination & Foliage Decoupling module with morphology-aware attention and a boundary-guided loss, explicitly tailoring the Transformer paradigm to the stringent demands of unstructured agricultural perception.

3. THE COMPLEX-FIELD-PEST (CFP) DATASET

To bridge the domain gap between laboratory-controlled phytopathology and real-world unstructured agricultural environments, we construct and introduce the Complex-Field-Pest (CFP) dataset. Unlike existing benchmarks that predominantly feature centered, isolated leaves against homogeneous backgrounds, CFP is meticulously curated to evaluate micro-lesion segmentation algorithms under severe, natural environmental stressors.

3.1 Data Collection and Environmental Diversity

The image acquisition process was executed across four major agricultural provinces over a continuous 12-month period, ensuring diverse phenological stages and weather conditions.

Hardware: Imagery was captured utilizing a multi-platform approach to ensure scale diversity. Canopy-level images were acquired using DJI Matrice 300 RTK Unmanned Aerial Vehicles (UAVs) equipped with Zenmuse P1 45-megapixel sensors. Ground-level macro photography was conducted using Sony A7R IV DSLR cameras equipped with 90mm macro lenses to capture extreme instances of micro-lesions and pest eggs.

Unstructured Challenges: We deliberately targeted unstructured scenarios. Of the 10,240 curated images, 35% contained strong specular reflections (caused by morning dew or waxy leaf cuticles under direct sunlight), 30% featured heavy physical occlusion (overlapping foliage and crossing branches), and 15% exhibited motion blur induced by wind dynamics during UAV capture. All raw 4K images were thoroughly filtered to remove out-of-focus artifacts, and uniformly downsampled to 2048×2048 for high-resolution annotation, before being cropped to 512×512 during network training.

Table 1 Target Size Distribution in the CFP Dataset.

Size Category	Pixel Area Range	Instance Count	Percentage (%)	Typical Pathological Examples
Extra-Small	≤ 5 px	46,320	42.5%	Early rust spores, mite eggs, initial fungal hyphae
Small	5 – 15 px	31,504	28.9%	Developing leaf spots, early blight lesions
Medium	15 – 50 px	18,966	17.4%	Mature bacterial spots, aphid colonies
Large	> 50 px	12,210	11.2%	Advanced necrosis, large miner tracks, coalesced lesions
Total	-	109,000	100%	-




Figure 1 Visual profiling of the CFP dataset: (a) Examples of unstructured field challenges: specular reflections (top), overlapping heavy occlusions (middle), and wind-induced motion blur (bottom); (b) Demonstration of the sub-pixel annotation protocol, where agricultural experts precisely delineate the fuzzy chlorotic halos of micro-lesions, avoiding the coarse approximations seen in standard datasets.

3.2 Sub-Pixel Annotation Protocol

A significant bottleneck in existing datasets is the reliance on coarse bounding boxes or rough polygonal masks, which inherently fail to capture the radial diffusion halos of biological diseases. To ensure that the ground-truth annotations aligned with our Level-Set optimization goals, we implemented a rigorous, three-stage sub-pixel annotation protocol:

- 1) **Initial Delineation:** Trained agricultural graduate students utilized interactive segmentation tools (e.g., LabelMe with SAM-assisted pre-labeling) to draft the initial contours.
- 2) **Expert Refinement:** Three senior agronomists independently reviewed and refined the contours at a 400% zoom level. They were specifically instructed to annotate the chlorotic transition zones (the fading yellow/brown halos surrounding the necrotic core) rather than just the central dead tissue.
- 3) **Consensus and Verification:** For ambiguous boundaries, a multi-expert voting mechanism was employed. The final masks were exported as high-precision semantic bitmasks, ensuring zero information loss regarding morphological edge irregularities.

3.3 Category Distribution and Scale Statistics

The CFP dataset comprised eight major economic crops (apple, corn, grape, citrus, and tomato) and covered 12 distinct categories of phytopathological targets (comprising fungal infections, bacterial blights, leaf miner tracks, and pest eggs).

As detailed in Table I, the defining characteristic of the CFP dataset is its extreme long-tail distribution of target sizes, which is a well-known bottleneck in computer vision [11], closely mirrors actual early-stage outbreak scenarios. Specifically, 42.5% of the annotated instances occupy an area of ≤ 5 pixels (defined as Extra-Small targets), making CFP significantly more challenging than standard benchmarks.



3.4 Unprecedented Complexity: Comparison with Existing Dataset

To unequivocally demonstrate the necessity of introducing the CFP dataset, Table 2 compares its core attributes with two widely utilized public benchmarks: PlantVillage[26] and RoCoLe[23].

While PlantVillage offers taxonomic diversity, it is severely limited by its laboratory-controlled, single-leaf backgrounds, rendering it unsuitable for evaluating field robustness. Conversely, while RoCoLe captures field conditions, it is constrained to a single crop (coffee) and possesses a relatively small sample size. The CFP dataset

Table 2 Comprehensive Comparison of CFP with Existing Agricultural Datasets.

Dataset Feature	PlantVillage	RoCoLe	CFP (Ours)
Total Images	54,303	1,560	10,240
Environment	Lab (Homogeneous)	Field (Unstructured)	Field (Highly Unstructured)
Crop Diversity	14 Species	1 Species (Coffee)	8 Species
Annotation Level	Coarse Polygon	Polygon	Sub-pixel Level-Set Contour
Micro-Lesions (≤ 5 px)	$\approx 15\%$	$\approx 25\%$	42.5%
Quantified Interferences*	None (Controlled)	Not Explicitly Quantified	65% (Occlusion/Specular)

**Figure 2** Statistical distribution of lesion sizes across different datasets. The CFP dataset uniquely emphasizes the “Extra-Small” (≤ 5 px) category, providing a rigorous benchmark for early-stage disease perception algorithms compared to the macroscopic bias of existing datasets.**Figure 3** The comprehensive architecture of AgriSpot-Trans. The input canopy image undergoes soft feature decoupling via the E-Decoupler. The 4-stage MRA Encoder extracts morphology-aware features, which are subsequently fused by the Level-Set Decoder to predict the SDF for sub-pixel boundary delineation.

bridges these critical gaps. It not only provides over 10,000 high-resolution field images of multiple crop species, but uniquely quantifies environmental interference (65% interference rate). Most crucially, the massive volume of micro-lesions (≤ 5 px) in CFP establishes it as the first benchmark explicitly optimized for evaluating early-stage, sub-pixel phytopathology segmentation.

4. METHODOLOGY

To robustly segment micro-lesions in unstructured agricultural environments, we propose AgriSpot-Trans, an end-to-end, occlusion-aware hierarchical Transformer. Unlike

conventional pipelines that crop input images based on hard masks—which inevitably disrupts global leaf topology—our architecture seamlessly integrates environmental decoupling, morphology-aware feature extraction, and continuous boundary evolution.

4.1 Overall Architecture of AgriSpot-Trans

The overall pipeline of AgriSpot-Trans is designed to process high-resolution agricultural imagery while maintaining sub-pixel precision for targets as small as $3 \sim 5$ pixels. As illustrated in Figure 3, the network processes an input image $I \in \mathbb{R}^{H \times W \times 3}$ via three tightly-coupled macro-modules:




Figure 4 Detailed schematic of the E-Decoupler. Instead of destructive hard-cropping, tokens representing optical noise are evaluated by an auxiliary MLP and suppressed by a learnable damping factor γ via a Gumbel-Softmax mechanism.

1. **Tokenization and Decoupling:** The image is partitioned into overlapping patches and fed into the Environment-Aware Token Routing (E-Decoupler) module. This performs soft-gating in the feature space, dampening tokens that correspond to severe optical interferences or physical occlusions, yielding a refined feature map.
2. **Morphology-Aware Feature Extraction:** A 4-stage hierarchical encoder powered by the Morphological Radial Attention (MRA) mechanism processes the purified tokens. It generates a multi-scale feature pyramid $\{Z_1, Z_2, Z_3, Z_4\}$ that explicitly models the radial diffusion halos of lesions.
3. **Level-Set Boundary Evolution:** A CNN-Transformer hybrid decoder fuses these features to predict a continuous Signed Distance Function (SDF). Guided by the Boundary-Aware Active Contour Loss (L_{BAAC}), the zero-level set dynamically evolves to encapsulate fuzzy micro-lesions.

4.2 Environment-Aware Token Routing (E-Decoupler)

In unstructured fields, cameras frequently capture waxy leaf reflections and dappled shadows. To suppress these optical noises without discarding spatial structure, we design the E-Decoupler as a differentiable soft-routing gate at the initial embedding layer (detailed in Figure 4).

Let the initial projection of the image patches be $X \in \mathbb{R}^{N_p \times C}$, where $N_p = \frac{H}{4} \times \frac{W}{4}$. An auxiliary Multi-Layer Perceptron (MLP) predicts an environmental confidence probability Π for each token:

$$\Pi = \sigma(W_2 \cdot \text{GELU}(W_1 \cdot \text{LN}(X) + b_1) + b_2) \in [0, 1]^{N_p}$$

To enable backpropagation through a discrete routing decision, we utilize the Gumbel-Softmax estimator[18]. The routing mask $M \in \{0, 1\}^{N_p}$ is sampled via:

$$M_i = \frac{\exp((\log(\Pi_i) + g_1)/\tau)}{\exp((\log(\Pi_i) + g_1)/\tau) + \exp((\log(1 - \Pi_i) + g_0)/\tau)}$$

where $g_0, g_1 \sim \text{Gumbel}(0, 1)$ are i.i.d. noise samples, and τ is the temperature parameter.

Instead of applying a hard mask ($X_i \cdot M_i$), we introduce a learnable semantic damping factor $\gamma \in [0, 1]$. The decoupled token sequence $\hat{X}$ is formulated as:

$$\hat{X}_i = (1 - \gamma(1 - M_i))X_i$$

Tokens identified as environmental noise ($M_i \approx 0$) are scaled down by $(1 - \gamma)$, preserving the absolute spatial grid while minimizing the impact of noise.

4.3 Morphological Radial Attention (MRA)

The core limitation of standard Vision Transformers in phytopathology is their reliance on Window-based Multi-head Self-Attention (W-MSA). Biological damages typically spread radially outward. As compared in Figure 5, a square window indiscriminately aggregates background healthy tissue at its corners, diluting the weak disease signal.

We propose Morphological Radial Attention (MRA). For a query token $q_{x,y}$, MRA casts N_{ray} continuous rays outward at uniform angular intervals $\theta_k = \frac{2\pi k}{N_{ray}}$. The physical coordinates of the sampled key/value points (x', y') are defined via polar-to-Cartesian mapping:

$$x' = x + r_j \cos(\theta_k), \quad y' = y + r_j \sin(\theta_k)$$

where r_j is the j -th sampling radius. The exact feature vectors for Keys (K_{rad}) and Values (V_{rad}) are obtained using bilinear interpolation.

The attention operation is then calculated using a specialized polar relative position bias B_{rad} :

$$\text{MRA}(Q, K_{rad}, V_{rad}) = \text{Softmax} \left(\frac{QK_{rad}^T}{\sqrt{d_k}} + B_{rad}(\Delta r, \Delta \theta) \right) V_{rad}$$

where $\Delta r = \sqrt{(x - x')^2 + (y - y')^2}$ and $\Delta \theta = \arctan 2(y' - y, x' - x)$. This explicitly guides the network to align its receptive field with the radial diffusion halos of micro-lesions.




Figure 5 Receptive field comparison. (a) Standard W-MSA incorporates irrelevant healthy tissue at the corners of its rigid square window. (b) Our proposed MRA employs a radial sampling strategy, dynamically aligning with the morphological spread of the biological lesion.




Figure 6 The 4-stage hierarchical MRA encoder. Patch merging layers build the multi-scale pyramid. Within each stage, successive blocks alternate between standard MRA and Shifted-MRA (SMRA) by rotating the sampling rays by $\frac{\pi}{N_{ray}}$ to ensure dense spatial coverage.

4.4 Hierarchical Encoder and Phase-Shifted MRA

To capture both the microscopic texture of early lesions and the macroscopic topology of the leaf, the encoder constructs a 4-stage feature pyramid, depicted in Figure 6.

Between each stage, a Patch Merging layer concatenates the features of 2×2 neighboring patches and projects the channel dimension to $2C$, systematically halving the spatial resolution from $\frac{H}{4} \times \frac{W}{4}$ down to $\frac{H}{32} \times \frac{W}{32}$.

To facilitate cross-ray information exchange and prevent isolated receptive fields, we introduce the Shifted-MRA (SMRA) mechanism. In successive Transformer blocks within the same stage, the angular sampling phase is rotated by $\Delta\phi = \frac{\pi}{N_{ray}}$. This alternating pattern (MRA $\rightarrow$ SMRA $\rightarrow$ MRA) ensures that the radial rays interleave, providing a dense, star-shaped coverage that maintains linear computational complexity $\mathcal{O}(HWC^2)$ while capturing global semantic context.

4.5 Level-Set Active Contour Decoder

1) Recovering sub-pixel boundaries requires transitioning from discrete pixel classification to continuous mathematical

contour modeling. We formulate our decoder $g_{decode}(\cdot)$ based on the Level-Set method (LSM)[16,25], illustrated in Figure 7.

A Feature Pyramid Network (FPN)[3] progressively up-samples and fuses $\{Z_1, Z_2, Z_3, Z_4\}$. The terminal layer is a 1×1 convolution without non-linear activation, which regresses a continuous Signed Distance Function (SDF), denoted as $\phi \in \mathbb{R}^{H \times W \times 1}$.

For any spatial location $x \in \Omega$, the SDF encodes the shortest Euclidean distance to the lesion boundary $\partial\Omega_{lesion}$. The contour topology is defined strictly as:

$$\begin{cases} \phi(x) > 0, & \text{if } x \text{ is inside the lesion} \\ \phi(x) = 0, & \text{if } x \text{ is on the exact boundary } \partial\Omega_{lesion} \\ \phi(x) < 0, & \text{if } x \text{ is in the background} \end{cases}$$

By extracting the zero-level set ($\phi(x) = 0$) via interpolation, AgriSpot-Trans escapes the constraints of the discrete camera grid.

4.6 Boundary-Aware Active Contour Loss Function (LBAAAC)

To optimize the continuous SDF ϕ , we propose a mathematically rigorous Boundary-Aware Active Contour Loss




Figure 7 Level-Set decoding process driven by L_{BAAC} . The decoder regresses a 3D Signed Distance Function (SDF) surface ϕ . The intersection with the zero-plane produces the continuous zero-level set (red contour), which is actively pulled towards the true lesion margin by L_{region} and smoothed by L_{length} .

(L_{BAAC}). We introduce smooth approximations for the Heaviside step function $H_\epsilon(\phi)$ and the Dirac delta function $\delta_\epsilon(\phi)$:

$$H_\epsilon(\phi) = \frac{1}{2} \left(1 + \frac{2}{\pi} \arctan \left(\frac{\phi}{\epsilon} \right) \right), \quad \delta_\epsilon(\phi) = \frac{\epsilon}{\pi(\epsilon^2 + \phi^2)}$$

The active contour optimization is driven by two PDE-inspired integral terms. The Region Fitting Term (L_{region}) calculates the Mean Squared Error against the ground-truth Y :

$$L_{region} = \int_{\Omega} (H_\epsilon(\phi(x)) - Y(x))^2 dx$$

The Length Penalty Term (L_{length}) computes the total curve length of the zero-level set, penalizing jagged boundaries caused by camera noise:

$$L_{length} = \int_{\Omega} \delta_\epsilon(\phi(x)) |\nabla \phi(x)| dx$$

To combat extreme class imbalance in early-stage disease detection, we integrate a Focal Loss (L_{focal}) [12]. The end-to-end training objective is:

$$L_{total} = \lambda_{AC} (L_{region} + \alpha L_{length}) + \lambda_{focal} L_{focal}$$

where $\lambda_{AC} = 1.0$, $\alpha = 0.1$, and $\lambda_{focal} = 0.5$ are empirically determined coefficients.

5. EXPERIMENTAL RESULTS

5.1 Experimental Setup

1) Datasets

To evaluate the performance of the proposed AgriSpot-Trans comprehensively, we conducted extensive experiments on three distinct datasets, progressing from controlled to highly unstructured field environments:

PlantVillage[26]: A widely adopted public benchmark containing 54,303 laboratory-controlled images spanning 14 crop species and 26 disease categories. Micro-lesions account for approximately 15% of the total annotated targets.

The dataset was split into 70% training, 10% validation, and 20% testing sets.

RoCoLe (Robusta Coffee Leaf)[23]: A public field dataset comprising 1,560 images capturing coffee leaf rust and mite damage under natural lighting. Extremely small spots (≤ 5 px) constitute around 25% of the annotations. It was partitioned using a standard 70/10/20 split.

Complex-Field-Pest (CFP) Dataset (Ours): To address the limitations of existing datasets, we constructed and introduced the CFP dataset, a large-scale, unstructured field benchmark. It contains 10,240 ultra-high-resolution images (captured at 4K resolution using UAVs and handheld smartphones, downsampled for training) covering 8 major crops and 12 types of micro-targets (e.g., early blight, rust spores, leaf miner tracks, and pest eggs). Crucially, 42% of the targets are extreme micro-lesions (≤ 5 px), and 65% of the images contain severe physical occlusions or optical interferences (e.g., specular reflections from morning dew, canopy shadows). The dataset was meticulously annotated by agricultural experts using polygon sub-pixel refinement tools, and divided strictly into 70% for training, 10% for validation, and 20% for testing.

2) Evaluation Metrics

Segmentation performance is evaluated using standard metric protocols: mean Intersection-over-Union (mIoU), Recall, Precision, and F1-score. To specifically validate our claims regarding micro-targets and morphological fuzzy boundaries, we introduce two specialized metrics: (1) Pixel-Binned Recall, which calculates recall across distinct target size thresholds (e.g., Extra-Small ≤ 5 px), and (2) Boundary F1-score (BF1), which strictly evaluates the contour alignment accuracy. Furthermore, to assess edge deployment viability, we report the number of Parameters (Params), floating-point operations (FLOPs), and Frames Per Second (FPS) measured on an NVIDIA Jetson Xavier NX edge computing device.

3) Implementation Details

The framework was implemented using PyTorch 2.0.1 and trained on a single NVIDIA RTX 3090 GPU. Input images were uniformly resized to 512×512 resolution with a batch size of 16. We optimized the network using the AdamW

optimizer with a weight decay of 0.01. The initial learning rate was set to 6×10^{-5} , governed by a cosine annealing learning rate scheduler following a 10-epoch linear warmup. The total training process spanned 200 epochs. Network weights were initialized using ImageNet pre-trained weights. To strengthen domain robustness, we applied agricultural-specific data augmentation, including Random ColorJitter, simulated synthetic dew reflections, and random shadow pasting. The hyperparameters for the L_{BAAC} loss function were strictly fixed as defined in our methodology: $\lambda_{AC} = 1.0$, $\alpha = 0.1$, and $\lambda_{focal} = 0.5$.

5.2 Comparison with State-of-the-Art Methods

We benchmarked AgriSpot-Trans against a comprehensive suite of semantic segmentation architectures, strictly categorized into classical CNNs, mainstream Vision Transformers (ViTs), and the latest SOTA models (including agricultural-specific frameworks).

As shown in Table I, AgriSpot-Trans achieves state-of-the-art results across all three datasets. Most notably, on the highly challenging CFP dataset, our model achieves a remarkable mIoU of 78.2%. This represents a 9.5% absolute improvement over the strong SegFormer baseline ($p < 0.01$), confirming the statistical significance of our architectural enhancements. Furthermore, our method outperforms the latest Mask2Former[9] and domain-specific Ag-SOTA models[15,20], excelling particularly in Precision and F1-score due to its suppression of optical false alarms.

5.3 Ablation Studies

To rigorously validate the independent contribution of each proposed module, we conducted comprehensive ablation studies on the CFP dataset. We used a pure SegFormer architecture as the absolute baseline to ensure single-variable isolation.

As demonstrated in Table II, the progressive integration of our modules yields a steady performance trajectory. The introduction of the E-Decoupler immediately increases mIoU by 2.7%, validating its capacity to filter background noise. Substituting standard W-MSA with our Morphological Radial Attention (MRA/SMRA) yields a massive 6.3% jump in ≤ 5 px Recall, proving that radial sampling effectively captures lesion diffusion halos. Finally, the inclusion of the SDF Decoder optimized by the active contour loss pushes the Boundary F1-score (BF1) to 81.5%, culminating in the full AgriSpot-Trans architecture.

5.4 Micro-Lesion & Sub-pixel Boundary Analysis

The defining challenge of early-stage phytopathology is the identification of extremely sparse targets. To validate our claims, we performed a pixel-binned recall analysis. As illustrated in Figure 11/Table III, while standard ViTs and

Ag-SOTA models maintain competitive accuracy on medium and large targets (> 15 px), their performance experiences a catastrophic drop when targeting extreme micro-lesions. In stark contrast, AgriSpot-Trans maintains a highly resilient recall curve. Specifically, in the Extra-Small (≤ 5 px) interval, our model achieves a recall of 76.0%, representing a massive 23.7% absolute improvement over the SegFormer baseline (52.3%).

This sub-pixel precision is driven directly by our Level-Set formulation. Figure 11 depicts the dynamic evolution of the active contour during training. By Epoch 50, the zero-level set perfectly maps the fuzzy, irregular biological margin, demonstrating that L_{BAAC} successfully prevents the network from collapsing into bounding-box-like predictions.

5.5 Robustness and Efficiency for Edge Deployment

To evaluate the deployment viability of agricultural drones and robotic rovers, we rigorously tested the model's environmental robustness and computational efficiency.

We partitioned the CFP test set into five extreme sub-scenarios. The radar chart in Figure 13 demonstrates that AgriSpot-Trans has no performance bottlenecks. Unlike Mask2Former, which drops significantly under Strong Illumination, our E-Decoupler ensures consistent high mIoU across all axes, definitely closing the domain gap between lab and field data.

Finally, Figure 14 illustrates the trade-off between deployment efficiency and precision. Deployed on an NVIDIA Jetson Xavier NX via TensorRT, AgriSpot-Trans resides firmly in the top-left quadrant of the FLOPs-Accuracy space (high accuracy, low computational load). Despite incorporating sub-pixel boundary evolution, the linear complexity of the MRA encoder allows the model to process 512×512 frames at 48 FPS (only 42.3 GFLOPs, 18.5M Params). This confirms that AgriSpot-Trans achieves a pristine balance: delivering SOTA pixel-level pathology perception while meeting the strict real-time, low-power constraints of precision agriculture edge hardware.

6. DISCUSSION

The extensive experimental results demonstrate that AgriSpot-Trans achieves state-of-the-art performance in segmenting agricultural micro-lesions under highly unstructured field conditions. Beyond the quantitative improvements, this section elucidates the underlying mechanisms driving these results, explores the practical implications for precision agriculture, and discusses current limitations to pave the way for future research.

6.1 Methodological Insights: Why AgriSpot-Trans Works

The superior performance of AgriSpot-Trans—particularly its 23.7% improvement in recall for targets ≤ 5 px—stems

Table 3 Quantitative Comparison with SOTA Methods across Three Agricultural Datasets. (Best results are in bold, second best are underlined. SegFormer is used as the baseline for improvement calculation.)

Method Category	Model	PlantVillage (mIoU) (%)	RoCoLe (mIoU) (%)	CFP (mIoU) (%)	CFP (Recall) (%)	CFP (Precision) (%)	CFP (F1-score) (%)
CNN-based	U-Net[17]	79.4	64.2	58.1	62.3	68.5	65.2
	DeepLabV3+[21]	82.1	67.5	61.4	65.8	70.2	67.9
	HRNet[14]	83.5	69.1	63.7	67.4	72.1	69.6
	SETR[5]	84.2	70.3	66.8	71.2	74.3	72.7
ViT-based	SegFormer (Baseline)[1]	86.8	73.4	68.7	73.5	76.8	75.1
	Swin-Unet[10]	87.1	74.5	69.4	75.1	77.2	76.1
	Mask2Former[9]	88.9	76.8	74.3	80.2	82.4	81.2
Latest/Ag SOTA	Ag-SOTA (AgriSeg)[15]	88.5	76.1	73.9	79.8	81.5	80.6
Proposed	AgriSpot-Trans	91.2	81.4	78.2 (+9.5%)	86.4	88.1	87.2




Figure 8 Qualitative visual comparison on the CFP dataset. Red dashed boxes indicate 2x/4x zoomed-in regions showcasing micro-lesions and fuzzy boundaries. Our method effectively eliminates false positives caused by specular reflections and precisely delineates sub-pixel biological boundaries.

Table 4 Ablation Study of Proposed Components on the CFP Dataset.

Config	E-Decoupler	MRA / SMRA	SDF Decoder & Loss	mIoU (%)	Recall (= 5 px) (%)	BF1 (%)
①	✗	✗	✗	68.7	52.3	64.1
②	✓	✗	✗	71.4 (+2.7%)	55.1	67.2
③	✓	✓	✗	75.3 (+3.9%)	61.4	72.8
④	✓	✓	✓	77.1 (+1.8%)	72.8	78.4
⑤(Full)	✓	✓	✓ L_{BAAC}	78.2 (+1.1%)	76.0	81.5




Figure 9 Interpretability of the E-Decoupler. The soft-routing mask actively suppresses weights for regions with specular reflections and deep shadows, delivering a purified feature map to the downstream MRA encoder.

Table 5 Pixel-Binned Recall Analysis on the CFP Dataset.

Size Category	SegFormer (Baseline)	Mask2Former	AgriSpot-Trans (Ours)
Large (> 50 px)	88.4%	91.2%	92.5%
Medium (15–50 px)	81.2%	85.6%	88.3%
Small (5–15 px)	67.5%	73.4%	82.1%
Extra-Small (≤ 5 px)	52.3%	61.8%	76.0% (+23.7%)

from a fundamental paradigm shift: moving away from generic computer vision pipelines towards a biologically and physically grounded architecture.

Firstly, the E-Decoupler addresses the optical realities of the field. Traditional methods that rely on hard-cropping or bounding boxes to isolate regions of interest inevitably

sever the continuous topological structure of the leaf canopy. By employing a Gumbel-Softmax-driven[18] soft-routing mechanism, our model actively dampens the feature weights of severe optical interferences (e.g., specular reflections from waxy leaves, dappled shadows) without destroying the spatial grid. This is analogous to the selective attention of the human




Figure 10 Grad-CAM feature heatmaps comparing standard W-MSA with our MRA. While W-MSA introduces background healthy tissue (square artifact), MRA perfectly wraps around the radial diffusion edge of the rust spot.




Figure 11 Pixel-binned recall analysis on the Complex-Field-Pest (CFP) dataset. Targets are categorized into four intervals based on pixel area. While standard ViTs and Mask2Former suffer a catastrophic performance degradation on extreme micro-lesions, AgriSpot-Trans demonstrates exceptional robustness, achieving a 76.0% recall rate on Extra-Small targets (≤ 5 px)—an absolute improvement of 23.7% over the baseline.




Figure 12 Dynamic evolution of the zero-level set (red contour) predicted by the SDF decoder across training epochs. The active contour progressively adapts to the irregular, sub-pixel diffusion margin of the lesion.

visual system, ensuring the downstream encoder receives a refined feature set.

Secondly, the Morphological Radial Attention (MRA) rectifies a critical morphological misalignment in standard Vision Transformers. Standard W-MSA employs rigid, square windows that inadvertently aggregate irrelevant healthy tissue at the corners when processing circular or irregular disease spots. In nature, biological damage (such as fungal sporulation or chlorotic degradation) spreads radially

outward from the epicenter of an infection. By utilizing a star-shaped sampling strategy parameterized by polar coordinates (r, θ) , MRA natively aligns its receptive field with this biological diffusion gradient, effectively isolating the weak pathological signals of early-stage lesions from the background chlorophyll.

Finally, the Level-Set SDF Decoder provides a mathematical escape from the discrete camera grid. Standard semantic segmentation models treat micro-lesion detection




Figure 13 Robustness evaluation across 5 extreme unstructured field scenarios. AgriSpot-Trans demonstrates superior resilience against environmental noise compared to baseline models.




Figure 14 mIoU vs. FLOPs bubble chart on Jetson Xavier NX. Bubble size represents model Parameters. AgriSpot-Trans occupies the optimal top-left quadrant, achieving the highest precision with edge-friendly computational demands.

as a binary pixel-classification problem, which forces the network to blindly guess the fuzzy boundaries of early-stage spots. By reformulating the output as a continuous Signed Distance Function $\phi(x)$ guided by the L_{BAAC} loss[16,25], the network dynamically evolves a sub-pixel contour. This continuous PDE-inspired approach is the direct catalyst for the model's unprecedented Boundary F1-score (81.5%) and its ability to effortlessly delineate targets smaller than the standard convolutional receptive field.

6.2 Implications for Precision Agriculture

The practical value of AgriSpot-Trans extends far beyond benchmark leaderboards. In the context of modern precision agriculture, early intervention is paramount. Achieving a 76.0% recall rate on extra-small disease spots allows agronomists to detect outbreaks days or even weeks before macroscopic defoliation occurs, saving both chemical costs and crop yields.

Furthermore, the high-fidelity, sub-pixel boundaries generated by our Level-Set decoder directly empower next-generation Variable Rate Application (VRA) systems. For agricultural UAVs equipped with pinpoint spraying nozzles, a generic bounding box is insufficient, as it leads to the

spraying of pesticides on surrounding healthy tissue and soil, causing ecological runoff. Our precise contours enable millimeter-level targeted spraying. Coupled with its demonstrated efficiency—running at 48 FPS on an edge-computing NVIDIA Jetson Xavier NX device—AgriSpot-Trans proves itself to be highly viable for real-time, closed-loop deployment on autonomous agri-bots operating in all-weather field environments.

6.3 Limitations of the Current Study

Despite its robust performance, we acknowledge certain limitations in the current framework. First, the model currently relies exclusively on monocular RGB imagery. When a single leaf is afflicted by concurrent, overlapping diseases with highly similar color profiles (e.g., early-stage rust conjoined with brown spot), the Level-Set boundary evolution occasionally merges them into a single pathological entity, failing to distinctly separate the specific disease types based solely on visible spectrum features.

Second, while the model maintains a lightweight computational footprint at a resolution of 512×512 , scaling the continuous SDF decoding process directly to native, unpadded 4K UAV imagery introduces a substantial GPU

memory bottleneck. The dense gradient computations required by the Length Penalty Term (L_{length}) of the active contour loss restrict ultra-high-resolution batch processing on standard edge hardware.

6.4 Future Work

To address these limitations and further advance field-based phytopathology, our future work will proceed in two main directions. First, we plan to extend the 2D spatial architecture into the temporal domain. By processing continuous video streams from agricultural rovers and incorporating optical flow algorithms, we aim to explicitly model and mitigate the severe motion blur caused by wind-induced canopy movement.

Second, we intend to improve the AgriSpot-Trans framework to accommodate multi-modal data fusion. By integrating Multispectral and Hyperspectral imaging capabilities, the network can leverage Near-Infrared (NIR) bands to perform chlorophyll fluorescence analysis. This will enable the model to peer beneath the visible leaf surface, effectively decoupling overlapping diseases based on their distinct biochemical and spectral signatures rather than relying solely on RGB morphology.

7. CONCLUSION

In this comprehensive study, we introduced AgriSpot-Trans, a novel, end-to-end occlusion-aware hierarchical Transformer specifically engineered to address the critical problem of pixel-level micro-lesion segmentation in highly unstructured agricultural environments. While existing semantic segmentation pipelines often struggle with the domain gap between laboratory-controlled imagery and complex field realities, our research proposed a fundamental paradigm shift. By adopting a “localize first, refine second” approach, we successfully transitioned from generic, discrete pixel classification to a biologically and physically grounded continuous boundary evolution framework.

The architectural contributions of AgriSpot-Trans are threefold, each explicitly designed to overcome distinct phytopathological challenges. First, the Environment-Aware Token Routing (E-Decoupler) mechanism effectively mitigates severe optical interferences, such as specular reflections and dappled shadows, by employing a Gumbel-Softmax-driven[18] soft-routing strategy that preserves the global topological integrity of the leaf canopy. Second, to address the morphological mismatch of standard square-window attention, we designed the Morphological Radial Attention (MRA). This star-shaped sampling mechanism natively aligns its receptive field with the irregular, radial diffusion gradients characteristic of biological tissue damage. Finally, we introduced a Level-Set Signed Distance Function (SDF) Decoder optimized by a novel Boundary-Aware Active Contour Loss (L_{BAAC}), allowing the network to escape the constraints of discrete camera grids and predict true sub-pixel boundaries for fuzzy disease margins.

Extensive empirical evaluations across three diverse benchmarks—PlantVillage, RoCoLe, and our newly curated

Complex-Field-Pest (CFP) dataset—unequivocally validate the superiority of the proposed framework. AgriSpot-Trans achieves a state-of-the-art mIoU of 78.2% in extreme field scenarios, outperforming the strong SegFormer baseline by a statistically significant margin of 9.5% ($p < 0.01$). Most remarkably, the model effectively solves the sparse micro-target detection problem, delivering an unprecedented 76.0% recall rate on extreme extra-small lesions (≤ 5 px). This represents a massive 23.7% absolute improvement over existing methodologies, achieved while maintaining a real-time computational efficiency of 48 FPS on edge-computing hardware.

Ultimately, the deployment of AgriSpot-Trans has significant implications for the future of precision agriculture. By enabling the highly accurate, early-stage detection of micro-lesions before macroscopic defoliation occurs, this framework provides the critical visual indicators necessary for next-generation Variable Rate Application (VRA) systems. Empowering agricultural UAVs and robotic rovers with such precise, sub-pixel perception will significantly reduce chemical pesticide usage, minimize ecological runoff, and ensure sustainable crop management in all-weather field environments. Future research will focus on extending this framework into the temporal domain and integrating multi-spectral data fusion to better distinguish complex, concurrent crop diseases.

REFERENCES

1. E. Xie, W. Wang, Z. Yu, A. Anandkumar, J. M. Alvarez, and P. Luo, “SegFormer: Simple and Efficient Design for Semantic Segmentation with Transformers,” arXiv:2105.15203 [cs], Oct. 2021, Available: <https://arxiv.org/abs/2105.15203>
2. Z. Liu et al., “Swin Transformer: Hierarchical Vision Transformer using Shifted Windows,” arXiv:2103.14030 [cs], Mar. 2021, Available: <https://arxiv.org/abs/2103.14030>
3. T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, “Feature Pyramid Networks for Object Detection,” arXiv:1612.03144 [cs], Apr. 2017, Available: <https://arxiv.org/abs/1612.03144>
4. T. Xiao, Y. Liu, B. Zhou, Y. Jiang, and J. Sun, “Unified Perceptual Parsing for Scene Understanding,” arXiv.org, Jul. 26, 2018. <https://arxiv.org/abs/1807.10221>
5. S. Zheng et al., “Rethinking Semantic Segmentation from a Sequence-to-Sequence Perspective with Transformers,” arXiv:2012.15840 [cs], Jul. 2021, Available: <https://arxiv.org/abs/2012.15840>
6. J. Acharya, Z. Sun, and H. Zhang, “Differentially Private Assouad, Fano, and Le Cam,” arXiv.org, 2020. <https://arxiv.org/abs/2004.06830> (accessed Aug. 18, 2025).
7. H. Zhang, V. Sindagi, and V. M. Patel, “Image De-raining Using a Conditional Generative Adversarial Network,” arXiv.org, 2017. <https://arxiv.org/abs/1701.05957> (accessed Aug. 18, 2025).
8. “Leaf Bidirectional Transmittance Distribution Function Estimates and Models for Select Deciduous Tree Species,” Ieee.org, 2017, doi: <https://doi.org/10.21227/yjek-2059>.
9. B. Cheng, A. G. Schwing, and A. Kirillov, “Per-Pixel Classification is Not All You Need for Semantic Segmentation,” arXiv:2107.06278 [cs], Oct. 2021, Available: <https://arxiv.org/abs/2107.06278>

10. H. Cao et al., “Swin-Unet: Unet-like Pure Transformer for Medical Image Segmentation”, arXiv:2105.05537 [cs, eess], May 2021, Available: <https://arxiv.org/abs/2105.05537>
11. M. Kisantat, Z. Wojna, J. Murawski, J. Naruniec, and K. Cho, “Augmentation for small object detection”, arXiv:1902.07296 [cs], Feb. 2019, Available: <https://arxiv.org/abs/1902.07296>
12. T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, “Focal Loss for Dense Object Detection”, arXiv:1708.02002 [cs], Feb. 2018, Available: <https://arxiv.org/abs/1708.02002>
13. F. Milletari, N. Navab, and S.-A. Ahmadi, “V-Net: Fully Convolutional Neural Networks for Volumetric Medical Image Segmentation”, arXiv:1606.04797 [cs], Jun. 2016, Available: <https://arxiv.org/abs/1606.04797>
14. J. Wang et al., “Deep High-Resolution Representation Learning for Visual Recognition,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 43, no. 10, pp. 3349–3364, 2021.
15. A. Ouamane et al., “Optimized Vision Transformers for Superior Plant Disease Detection,” *IEEE Access*, vol. 13, pp. 48552–48570, 2025, doi: 10.1109/ACCESS.2025.3547416.
16. D. Marcos, D. Tuia, B. Kellenberger, L. Zhang, M. Bai, R. Liao, and R. Urtasun, “Learning deep structured active contours end-to-end,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 8877–8885.
17. S. H. M. Ali et al., “Optimized Vision Transformers for Superior Plant Disease Detection,” *IEEE Xplore*, 2025.
18. E. Jang, S. Gu, and B. Poole, “Categorical reparameterization with Gumbel-Softmax,” in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2017.
19. J. Sun, Z. Shen, Y. Wang, H. Bao, and X. Zhou, “LoFTR: Detector-Free Local Feature Matching with Transformers”, arXiv:2104.00680 [cs], Apr. 2021, Available: <https://arxiv.org/abs/2104.00680>
20. Z. Wang et al., “LKCAFormer: a lightweight transformer with large-kernel cooperative attention for the segmentation of field maize leaf diseases,” *Plant Methods*, 2024. (via PubMed Central).
21. O. Ronneberger, P. Fischer, and T. Brox, “U-Net: Convolutional Networks for Biomedical Image Segmentation,” in *MICCAI*, 2015, pp. 234–241.
22. C. Camilo, “On the discontinuities of Hausdorff dimension in generic dynamical Lagrange spectrum”, arXiv.org, 2024. <https://arxiv.org/abs/2403.18940> (accessed Aug. 18, 2025).
23. J. Parraga-Alava, K. Cusme, A. Loor, and E. Santander, “RoCoLe: A robusta coffee leaf images dataset for evaluation of machine learning based applications in plant diseases recognition,” *Data in Brief*, vol. 25, p. 104414, 2019.
24. A. Kamilaris and F. X. Prenafeta-Boldú, “Deep learning in agriculture: A survey,” *Computers and Electronics in Agriculture*, vol. 147, pp. 70–90, 2018.
25. S. Osher and J. A. Sethian, “Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations,” *Journal of Computational Physics*, vol. 79, no. 1, pp. 12–49, 1988.
26. M. H. Saleem, J. Potgieter, and K. M. Arif, “Plant disease detection and classification by deep learning,” *Plants*, vol. 8, no. 11, p. 468, 2019.
27. R. Qian, R. T. Tan, W. Yang, J. Su, and J. Liu, “Attentive Generative Adversarial Network for Raindrop Removal from a Single Image”, arXiv.org, 2017. <https://arxiv.org/abs/1711.10098> (accessed Aug. 18, 2025).
28. Y. Lu, S. Yi, N. Zeng, Y. Liu, and Y. Zhang, “Identification of rice diseases using deep convolutional neural networks,” *Neurocomputing*, vol. 267, pp. 378–384, 2017.
29. X. Zhang, “Deep learning model for image classification in machine vision,” *Engineering Intelligent Systems*, vol. 33, no. 4, 2025.
30. S. Krinidis, M. Krinidis, and V. Chatzis, “Fast and robust fuzzy active contours,” *Engineering Intelligent Systems*, vol. 21, no. 2–3, 2013.
31. L. Guo, B. Luo, G. Zhang, J. Wu, B. Tu, and X. Ou, “Edge detection of infrared image base on phase congruency,” *Engineering Intelligent Systems*, vol. 24, no. 3–4, 2016.

