



Article

Weakly Supervised Leaf Disease Segmentation via Reinforcement Learning-Based Prompt Optimization for SAM

Jian Zheng ^{1,2}, Junfan Jin ^{2,*}, Yong Chen ², Xiangxing Zhou ², Wentao Huang ² and Shengxiong Huang ²

¹ Yichun Lithium New Energy Industry Research Institute, Jiangxi University of Science and Technology, No. 5 Chunhua Road, Yichun Economic Development Zone, Yichun 336000, China; 9120030026@jxust.edu.cn

² Jiangxi Provincial Key Laboratory of Multidimensional Intelligent Perception and Control, School of Information Engineering, Jiangxi University of Science and Technology, Ganzhou 341000, China; 6720240947@mail.jxust.edu.cn (Y.C.); 6720240870@mail.jxust.edu.cn (X.Z.); 6120240825@mail.jxust.edu.cn (W.H.); 6720240919@mail.jxust.edu.cn (S.H.)

* Correspondence: 6120240817@mail.jxust.edu.cn

Abstract

Accurate segmentation of leaf lesions is essential for crop disease monitoring and precision agriculture. However, most existing methods rely on pixel-level annotations, which are costly and difficult to obtain in real agricultural scenarios. Weakly supervised approaches alleviate this requirement but often suffer from incomplete lesion localization and background interference. In addition, recent foundation segmentation models demonstrate strong generalization ability, yet their performance heavily depends on prompt quality, which is difficult to design automatically. To address these challenges, this paper proposes a weakly supervised lesion segmentation framework that integrates probabilistic feature disentanglement and reinforcement learning-based prompt optimization. First, a probabilistic class activation mapping (CAM) is introduced to model lesion confidence and guide feature disentanglement, effectively suppressing background noise. Second, a dual-loss constraint is designed to improve pseudo-mask quality by enforcing both region consistency and structural regularity. Finally, a prompt optimization strategy is developed to automatically select informative prompts, enabling a foundation segmentation model to refine lesion boundaries without human interaction. Experimental results on multiple agricultural datasets demonstrate that the proposed method achieves superior performance compared with existing weakly supervised baselines. Specifically, the method achieves an mIoU of 45.6% and Dice score of 57.1% on apple leaf disease datasets. The results indicate that the proposed framework can effectively reduce annotation cost while providing reliable segmentation performance, indicating potential for practical agricultural applications.

Keywords: SAM; weakly supervised; disease segmentation; reinforcement learning; leaf lesion; probabilistic CAM; prompt optimization



Academic Editor: Ray E. Sheriff

Received: 12 August 2026

Revised: 12 October 2026

Accepted: 12 October 2026

Published: 15 October 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and

conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Plant diseases significantly affect crop yield and quality, posing a major challenge to sustainable agricultural production [1]. Among various disease phenotypes, leaf lesions are one of the most common and informative indicators for disease diagnosis and severity assessment. Accurate segmentation of lesion regions is therefore essential for automated crop monitoring and precision agriculture [2,3].

In recent years, deep learning-based semantic segmentation methods have achieved promising results in plant disease analysis [4,5]. Models such as U-Net [6], DeepLab [7]

and Transformer-based [8] architectures can provide accurate pixel-level predictions [9]. However, these approaches typically rely on large-scale pixel-level annotations, which are labor-intensive and require expert knowledge. In practical agricultural scenarios, collecting such annotations at scale is often infeasible due to time and cost constraints [10].

To alleviate this limitation, weakly supervised learning has been explored as an alternative paradigm, where models are trained using only image-level labels [11,12]. A common approach is to employ class activation mapping (CAM) [13] to localize discriminative regions. Despite its effectiveness, CAM-based methods often highlight only partial lesion regions and are sensitive to background interference, leading to incomplete or noisy pseudo masks [14]. This limitation is particularly pronounced in agricultural images, where complex backgrounds, leaf textures, and illumination variations can easily confuse the model.

Meanwhile, foundation segmentation models have recently demonstrated strong generalization capabilities across diverse visual tasks [15,16]. These models can produce high-quality segmentation results when guided by appropriate prompts. However, their performance heavily depends on the quality of the prompts. In agricultural applications, manual prompt design is impractical for large-scale deployment, and existing automatic prompt generation strategies are often heuristic and lack adaptability.

Therefore, three key challenges remain in weakly supervised leaf disease segmentation:

- (1) Unreliable lesion localization due to background interference;
- (2) Insufficient constraints in pseudo-mask generation under weak supervision;
- (3) Lack of effective automatic prompt optimization for foundation models.

The objective of this study is to develop a weakly supervised framework for accurate leaf disease lesion segmentation using only image-level supervision. Specifically, we aim to improve lesion localization, stabilize pseudo-mask generation, and automatically optimize prompts for a frozen foundation segmentation model. The proposed method consists of three stages: leaf extraction, weakly supervised lesion localization, and prompt-based segmentation refinement. Specifically, a probabilistic CAM-guided feature disentanglement strategy is introduced to suppress background interference and improve pseudo-mask quality. A dual-loss constraint is further designed to enhance the stability and structural consistency of lesion localization. In addition, an automatic prompt optimization strategy is developed to guide a foundation segmentation model for accurate lesion refinement without human intervention.

The main contributions of this work are summarized as follows:

- A probabilistic CAM-guided feature disentanglement strategy is proposed to improve lesion localization under weak supervision.
- A dual-loss constraint is designed to enhance pseudo-mask quality by enforcing region consistency and structural regularity.
- An automatic prompt optimization strategy is developed to improve segmentation performance without manual interaction.
- Experimental results demonstrate the effectiveness and robustness of the proposed method, while showing partial transferability across different crop scenarios.

2. Related Work

2.1. Fully Supervised Plant Disease Segmentation

Semantic segmentation has been widely studied in computer vision and has achieved significant success in plant disease analysis [17,18]. Convolutional neural network (CNN)-based models, such as U-Net and DeepLab, have been extensively applied for pixel-level segmentation of leaf lesions due to their strong feature extraction capability [19,20]. U-Net

adopts an encoder–decoder architecture with skip connections, enabling effective fusion of low-level and high-level features, and has been widely used in plant phenotyping tasks [6]. DeepLab further improves segmentation performance by introducing atrous convolution and spatial pyramid pooling, allowing multi-scale context modeling [21,22].

Researchers have made substantial improvements based on the above classic models. Shafiq and Gu [23] proposed MU-Net, which optimizes U-Net by introducing residual modules and residual paths, improving segmentation accuracy but remaining sensitive to background noise. Zeng and He [3] proposed an improved DeepLabv3+ network using ResNet101 as the backbone, incorporating a feature fusion branch based on a feature pyramid network in the encoder to achieve multi-level feature map fusion for black rot spot segmentation on grape leaves. Furthermore, methods based on Transformer have gradually emerged, such as UPFormer [24], SegFormer [25], and AISOA-SSformer [26]. UPFormer adopts a U-shaped hierarchical structure and generates small tokens to optimize computational cost. In addition, it combines Prototype Perception and Pixel-Aware Broadcasting to extract low-frequency global and high-frequency local features, which is crucial for accurately identifying small disease regions. AISOA-SSformer integrates a sparse global update perceptron and a saliency feature attention mechanism, achieving excellent performance in rice leaf disease segmentation. SegFormer achieves competitive performance while maintaining computational efficiency. These methods have achieved good results in grape and rice leaf disease segmentation, respectively. The emergence of the Mamba architecture provides a new solution. As a structure based on the State Space Model, Mamba [27] can more naturally maintain local continuity of sequences while efficiently capturing long-range dependencies by introducing time-varying parameters and hardware-aware algorithms [28]. Although these methods achieve high accuracy, they generally rely on a large amount of manual annotation and are designed for specific crops, resulting in limited generalization ability.

2.2. Weakly Supervised and Unsupervised Segmentation

To reduce annotation costs, weakly supervised and unsupervised methods have become research hotspots. Researchers explored weakly supervised segmentation methods based on image-level labels and proposed ResNet-CAM and WSLSS [29]. Although their accuracy is lower than that of fully supervised methods, they exhibit good generalization ability on unknown disease types. Xu et al. [30] proposed C²AM, which generates class-agnostic activation maps via contrastive learning, achieving unsupervised class-agnostic object segmentation. Subsequent methods such as WSCUOD [31], A2S-V2 [32], and FOUND [33] performed poorly on leaf disease datasets because they can only segment leaf regions but cannot effectively localize disease spots. Tian et al. [16] further proposed the UPLS framework, combining contrastive learning and the SAM [34] model to achieve annotation-free segmentation of leaf disease regions, and improved segmentation accuracy through a high-frequency attention module and multi-scale feature extraction. These methods provide feasible paths for reducing annotation dependence.

2.3. Foundation Model-Based Segmentation and Prompt Optimization

Foundation segmentation models have recently emerged as a powerful tool for general-purpose image segmentation [35]. Models such as the Segment Anything Model can generate high-quality segmentation masks when provided with appropriate prompts, demonstrating strong generalization across diverse visual domains.

To apply these models in automated settings, several studies have attempted to design prompt generation strategies. Existing approaches include random sampling, geometric heuristics, and feature similarity-based methods to generate point or box prompts. Several

studies have attempted to achieve automatic prompt engineering for SAM by generating geometric prompts [36]. Liu et al. [37] computed the feature similarity between pixels in a target image and a reference image, creating a point prompt scheme using only a single reference image. Liu et al. [38] introduced Matcher, an instance-level bidirectional matching strategy for prompt point generation. Kumari and Bansal [39] also randomly selected prompt points from masks produced by coarse segmentation, fine-tuned them, and then fed them into SAM to obtain segmentation results. The quality of the prompts generated by these methods is often limited because they do not fully account for the relationship between positive and negative prompts, leading to redundant and erroneous prompts that negatively affect segmentation performance.

More recently, reinforcement learning-based approaches have been explored to improve prompt selection by modeling it as a decision-making problem. The introduction of PPO [11] enables the combination of SAM with reinforcement learning, turning it into a plug-and-play tool that directly improves model segmentation performance. While these methods show potential, they often depend heavily on the quality of initial pseudo masks and introduce additional computational complexity.

3. Methodology

3.1. Overview

To address the challenges of high annotation cost, complex lesion patterns, and background interference in agricultural images, this study proposes a three-stage weakly supervised framework for leaf lesion segmentation.

As illustrated in Figure 1, the framework consists of the following:

- (1) Leaf extraction to suppress background interference;
- (2) Weakly supervised lesion localization to generate reliable pseudo masks;
- (3) Prompt-based refinement to improve segmentation accuracy using a foundation model.

The core idea is to progressively improve pseudo-mask quality and leverage automatic prompt optimization for accurate lesion boundary refinement without human interaction.




Figure 1. Overall architecture of the model.

3.2. Weakly Supervised Lesion Localization

The goal of this stage is to generate accurate and stable pseudo masks using only image-level labels. To address the limitations of conventional CAM-based methods, we introduce two key components: probabilistic feature disentanglement and a dual-loss constraint.

3.2.1. Probabilistic CAM-Guided Feature Disentanglement (PCFD)

Conventional class activation mapping (CAM) often produces sparse and incomplete responses and is sensitive to background noise, especially in agricultural images with complex textures.

To improve localization reliability, we reformulate CAM as a probabilistic representation. Given a feature map F , the activation map is normalized to obtain a confidence map :

$$P(x, y) = \sigma(A(x, y)) \quad (1)$$

where $P(x, y) \in [0, 1]$ represents the probability of each pixel belonging to lesion regions. Compared with conventional CAM, this probabilistic formulation provides a smoother and more consistent estimation of lesion regions. Based on this representation, we perform feature disentanglement:

$$F^{fg} = P \odot F, \quad F^{bg} = (1 - P) \odot F \quad (2)$$

where F^{fg} and F^{bg} denote foreground and background features, respectively. This soft separation mechanism avoids the instability caused by hard thresholding and enables the model to suppress background interference while preserving uncertain lesion regions. As a result, more reliable pseudo masks can be generated for subsequent refinement.

3.2.2. Dual-Loss Constraint (DLC)

Under weak supervision, pseudo-mask generation is often unstable due to insufficient constraints. To address this issue, we introduce a dual-loss constraint that jointly enforces region consistency and structural regularity.

First, a region consistency loss is defined to align the predicted lesion regions with the probabilistic CAM prior:

$$\mathcal{L}_{reg} = \|M - P\| \quad (3)$$

where M denotes the predicted mask. Second, a structural constraint is introduced to regularize lesion shapes:

$$\mathcal{L}_{str} = \frac{\text{Perimeter}(M)}{\text{Area}(M)} \quad (4)$$

This term acts as a generic shape regularizer that discourages highly fragmented or excessively irregular predicted regions. It is not intended to assume a disease-specific lesion shape, but rather to provide a weak morphological prior that favors spatially coherent lesion regions.

$$\mathcal{L} = \mathcal{L}_{reg} + \lambda \mathcal{L}_{str} \quad (5)$$

where λ balances the two terms. By combining semantic consistency and structural constraints, the proposed loss improves the stability and reliability of pseudo-mask generation. The details are shown in Figure 2.




Figure 2. Model details.

To examine the potential instability caused by the perimeter-to-area structural regularization during early training, we monitored the training losses in the initial epochs. As shown in Figure 3, the total training loss decreased consistently from 2.4352 at epoch 0 to 2.0537 at epoch 8, while the major loss components also decreased without noticeable oscillation or divergence. These observations suggest that the structural regularization does not introduce apparent optimization instability in the early training stage under our experimental setting.




Figure 3. Loss.

To further examine the sensitivity of the structural regularization to boundary noise, we compared the proposed perimeter-to-area formulation with two commonly used shape descriptors in Table 1. Specifically, we evaluated the scale-normalized compactness

$$C_{PP} = \frac{4\pi A}{p^2}, \quad (6)$$

and the solidity

$$S = \frac{A}{A_{hull}}, \quad (7)$$

where A_{hull} denotes the area of the convex hull. Their corresponding regularization terms were defined as $1 - C_{\text{PP}}$ and $1 - S$, respectively. These comparisons were conducted under the same training configuration to evaluate the robustness of different morphological constraints.

Table 1. Comparison of different morphological regularization strategies.

Structural Regularization	mIoU (%)	Dice (%)
None	41.6	52.9
Perimeter-to-Area (P/A)	44.1	55.8
Normalized Compactness ($1 - \frac{4\pi A}{P^2}$)	43.7	55.2
Solidity ($1 - \frac{A}{A_{\text{hull}}}$)	43.5	54.9

3.3. Prompt-Based Segmentation Refinement

Although the pseudo masks provide coarse localization, they often lack precise boundaries. To address this issue, we refine segmentation results using a foundation segmentation model.

Given a pseudo mask M_p , candidate prompts are generated from lesion regions and surrounding areas. These prompts are used to guide a frozen segmentation model to produce refined masks:

$$M = \text{SAM}(I, \mathcal{P}) \quad (8)$$

where $\mathcal{P}$ denotes the selected prompt set. The key challenge is how to select informative prompts automatically, which is addressed by the proposed optimization strategy.

3.4. Automatic Prompt Optimization Strategy (RLPO)

To improve prompt selection, we formulate it as a decision-making problem and introduce an automatic prompt optimization strategy. See Figure 4 for RLPO.




Figure 4. RLPO details.

3.4.1. Problem Formulation

Given a pseudo mask M_p and a set of candidate prompts $\mathcal{C} = \{c_i\}_{i=1}^N$, the goal is to select a subset of prompts $\mathcal{P} \subset \mathcal{C}$ that maximizes the quality of the refined segmentation.

The process is modeled as a Markov decision process (MDP):

- State s_t : representation of the current prompt set and pseudo-mask context;
- Action a_t : selecting a candidate prompt c_i ;
- Policy $\pi_\theta(a_t|s_t)$: probability of selecting action a_t ;
- Reward R_t : segmentation improvement after applying the selected prompts.

3.4.2. Candidate Prompt Generation

The quality of the candidate prompt space directly affects the effectiveness of prompt optimization. Therefore, candidate points are generated from the lesion localization results rather than sampled independently from the entire image.

For each candidate point c_i , the policy receives an 8-dimensional basic feature vector:

$$f_i = [r_i, g_i, b_i, x_i, y_i, \mu_i, \sigma_i, m_i], \quad (9)$$

where r_i, g_i , and b_i denote the RGB values at the candidate location, x_i and y_i represent the normalized spatial coordinates, μ_i and σ_i describe local image appearance statistics, and m_i indicates whether the candidate point lies inside the initial pseudo-mask.

In addition to these basic appearance and spatial features, contextual information is incorporated to provide the policy with region-level cues. The additional feature vector is defined as

$$e_i = [d_i^{box}, b_i^{neg}, v_i], \quad (10)$$

where d_i^{box} represents the spatial relationship between the candidate point and the lesion-related region, b_i^{neg} indicates whether the point belongs to a negative candidate region, and v_i represents local appearance variation.

The complete candidate representation is therefore

$$z_i = [f_i, e_i]. \quad (11)$$

The basic feature branch and the additional contextual feature branch are embedded separately and subsequently fused by the policy network. This design enables the policy to jointly exploit point-level appearance, spatial information, pseudo-mask information, and region-level contextual cues.

During training, candidate points are sampled without replacement from the candidate distribution. This prevents the policy from selecting the same point repeatedly within a single prompt-selection episode.

3.4.3. Weakly Supervised Reward Design

A major challenge in weakly supervised prompt optimization is the absence of pixel-level ground-truth masks during policy learning. Therefore, the reward cannot rely exclusively on the ground-truth segmentation quality. We instead construct a weakly supervised proxy reward using the Stage-2 pseudo-mask as a reference.

Let M_p denote the initial pseudo-mask and M_t denote the mask produced by SAM after the current prompt selection. The proxy reward is defined as

$$R_{\text{proxy}} = \tanh(R_{\text{IoU}} + 0.3R_{\text{boundary}} - 0.5R_{\text{area}}) \quad (12)$$

Region consistency

The first term measures the consistency between the refined mask and the initial pseudo-mask:

$$R_{IoU} = IoU(M_t, M_p). \quad (13)$$

This term encourages the selected prompts to preserve the lesion regions identified by the weakly supervised localization stage.

Boundary response

To characterize the change in boundary response after SAM refinement, a boundary-energy term is calculated for the predicted mask. Let $B(M)$ denote the boundary energy of mask M . The boundary-response component is defined according to the increase relative to the initial pseudo-mask:

$$R_{\text{boundary}} = \frac{\max(0, B(M_t) - B(M_p))}{\sqrt{H^2 + W^2} \times 10}, \quad (14)$$

where H and W denote the height and width of the image, respectively.

This term provides the policy with information about changes in boundary response after prompt selection.

Area expansion constraint

Because inappropriate prompts may cause SAM to expand the lesion region excessively, we introduce an area-expansion penalty:

$$R_{\text{area}} = \max\left(0, \frac{|M_t|}{|M_p| + \epsilon} - 1.5\right), \quad (15)$$

where $|M|$ denotes the number of foreground pixels in mask M , and ϵ is a small constant for numerical stability.

The penalty is activated when the refined mask becomes substantially larger than the initial pseudo-mask. This constraint is particularly relevant to leaf disease segmentation because inappropriate prompts may cause SAM to include large non-lesion regions.

The combination of Equations (12)–(15) allows the reward to simultaneously consider lesion-region consistency, changes in boundary response, and excessive mask expansion. Importantly, this proxy reward can be calculated without pixel-level ground-truth annotations.

3.4.4. Policy Network and Policy-Gradient Optimization

The policy network maps the candidate features to a selection probability. The 8-dimensional basic feature vector is first projected into a 64-dimensional representation, while the three additional contextual features are separately projected into a lower-dimensional representation. The two representations are subsequently fused to obtain the final candidate embedding.

For candidate c_i , the policy network produces a scalar logit:

$$q_i = g_\theta(f_i, e_i). \quad (16)$$

The selection probability is obtained using a softmax function over the currently available candidate set:

$$\pi_\theta(a_i | s_t) = \frac{\exp(q_i)}{\sum_{j \in \mathcal{C}_t} \exp(q_j)}, \quad (17)$$

where $\mathcal{C}_t$ denotes the candidate points that remain available at step t .

Candidate points are selected sequentially without replacement. After a point is selected, it is removed from the candidate set before the next action is sampled.

The policy parameters are optimized using a policy-gradient objective. The policy-gradient loss for a selected prompt set is written as

$$L_{PG} = -\hat{R}_t \sum_{i \in \mathcal{P}_t} \log \pi_{\theta}(a_i | s_t), \quad (18)$$

where $\hat{R}_t$ denotes the normalized reward.

To encourage exploration during policy learning, an entropy regularization term is incorporated:

$$L_{RL} = L_{PG} - \lambda_H H(\pi_{\theta}), \quad (19)$$

where

$$H(\pi_{\theta}) = -\sum_i \pi_{\theta}(a_i | s_t) \log \pi_{\theta}(a_i | s_t) \quad (20)$$

is the policy entropy, and λ_H controls the entropy regularization strength. In the implementation, the entropy coefficient is set to 0.02.

The policy is therefore encouraged to assign higher probability to prompt selections that lead to higher rewards while maintaining sufficient exploration during training.

3.4.5. Inference Strategy

During inference, the learned policy is used deterministically to rank candidate prompts. The top- k prompts with the highest selection probabilities are chosen:

$$\mathcal{P}^* = \text{Top-}k(\pi_{\theta}) \quad (21)$$

This ensures stable and efficient segmentation without stochastic sampling. The detailed algorithm is shown in Algorithm 1.

Algorithm 1 RL-based Prompt Optimization (RLPO)

Require: Image I , pseudo mask M_p , candidate set $\mathcal{C}$

Ensure: Refined mask M

```

1: Initialize policy  $\pi_{\theta}$ 
2: Initialize prompt set  $\mathcal{P} \leftarrow \emptyset$ 
3: for  $t = 1$  to  $T$  do
4:   Construct state  $s_t$ 
5:   Compute  $\pi_{\theta}(a | s_t)$ 
6:   Select prompt  $a_t$ 
7:    $\mathcal{P} \leftarrow \mathcal{P} \cup \{a_t\}$ 
8:    $M_t \leftarrow \text{SAM}(I, \mathcal{P})$ 
9:    $R_t \leftarrow \text{IoU}(M_t, M_p)$ 
10:  Update  $\theta$ 
11: end for
12: Select top- $k$  prompts
13:  $M \leftarrow \text{SAM}(I, \mathcal{P})$ 
   return  $M$ 

```

3.4.6. Difference from Existing Prompt Optimization Methods

Existing automatic prompt-generation approaches for SAM commonly rely on random sampling, geometric heuristics, feature matching, or predefined rules. More recently, reinforcement learning has also been explored for automatic prompt optimization. For example, Plug-and-Play PPO formulates SAM point-prompt optimization as a reinforcement learning problem and uses PPO to optimize prompt selection.

The proposed RLPO differs from these approaches in its formulation for weakly supervised leaf disease segmentation. First, the candidate prompt space is constructed from the lesion localization results generated by the preceding weakly supervised stage, rather than being independent of the lesion pseudo-mask. Second, candidate points are represented using both appearance and spatial information together with pseudo-mask and region-context cues. Third, the policy is trained using a weakly supervised reward that evaluates consistency with the generated pseudo-mask while constraining excessive mask expansion. Finally, the SAM parameters remain frozen throughout prompt optimization, and only the prompt-selection policy is learned.

Therefore, the contribution of RLPO is not the introduction of reinforcement learning itself. Instead, RLPO provides a task-specific formulation that connects weakly supervised lesion localization with automatic point-prompt selection for a frozen foundation segmentation model.

4. Experiments

4.1. Datasets

To verify the effectiveness of the proposed method in agricultural disease segmentation tasks, this paper conducts experiments on a publicly available apple leaf disease dataset [40]. This dataset includes apple leaf images affected by five common diseases: blotch, apple scab, black rot, cedar apple rust, and rust. These images were collected under both laboratory-controlled and real-world agricultural conditions. Among the complete apple disease dataset of 11,270 images, 1089 images were collected under real-world agricultural conditions, including 869 *Alternaria* leaf spot images and 220 rust images. The remaining 10,181 images were obtained under laboratory conditions. Additionally, the dataset also includes grape and strawberry leaf disease datasets: these contain images of grape and strawberry leaf diseases and serve as additional test datasets to evaluate generalization to unseen leaf disease types. Detailed information is presented in Table 2.

Table 2. Dataset of apple, grape, and strawberry leaf diseases with pixel-level annotations.

Category	Disease Name	Source	Images
Apple diseased leaves	Apple scab	Plant Village	2028
	Black rot	Plant Village	2119
	Cedar apple rust	Plant Village	1247
	<i>Alternaria</i> leaf spot	Apple Leaf Dataset	1820
	Rust	Apple Leaf Dataset	4056
Grape & strawberry dataset	Grape black rot	Plant Village	1300
	Grape esca	Plant Village	1300
	Strawberry leaf blight	Plant Village	1209

4.2. Evaluation Metrics

We adopt standard segmentation metrics to quantitatively evaluate performance:

$$\text{mIoU} = \frac{1}{N} \sum_{i=1}^N \frac{|M_i \cap G_i|}{|M_i \cup G_i|} \quad (22)$$

$$\text{Dice} = \frac{2|M \cap G|}{|M| + |G|} \quad (23)$$

$$\text{Precision} = \frac{TP}{TP + FP} \quad (24)$$

$$\text{Recall} = \frac{TP}{TP + FN} \quad (25)$$

where TP , FP , and FN denote true positives, false positives, and false negatives, respectively. These metrics comprehensively measure segmentation accuracy, region overlap, and detection completeness.

4.3. Implementation Details

The framework is implemented in PyTorch 1.12.1. The backbone is initialized with pre-trained weights and fine-tuned using image-level disease labels only. During training, images are resized to 512×512 and randomly horizontally flipped. ImageNet mean and standard deviation are used for normalization. The batch size is 32, the initial learning rate is 0.001, the weight decay is 1×10^{-4} , and the default training schedule consists of 10 epochs. A polynomial learning-rate schedule is used during optimization.

For CAM-based inference, multi-scale evaluation is performed at scales of 0.5, 1.0, 1.5, and 2.0. At each scale, both the original and horizontally flipped images are evaluated, and the flipped CAM is mapped back before fusion. The resulting CAMs are summed across scales and normalized by their maximum activation. Unless otherwise specified, a threshold of 0.25 is used to obtain the binary localization mask. No CRF refinement is applied in the reported default setting (CRF iterations = 0). For prompt optimization, the number of candidate points is set to 50, and the number of selected prompts during inference is fixed to $k = 3$.

To assess the robustness of the experimental results, all trainable models and stochastic prompt-selection procedures were evaluated using five independent random seeds (0, 1, 2, 3, and 4). The same training and test partitions were maintained across all runs. For each metric, we report the mean and standard deviation across the five runs. The proposed method was evaluated across different disease categories and imaging conditions, with the results summarized in Table 3.

Table 3. Segmentation performance under field and laboratory imaging conditions for different apple leaf disease categories.

Disease	Condition	mIoU (%)	Precision (%)	Recall (%)
Apple scab	Laboratory	46.8	96.1	64.2
Black rot	Laboratory	47.2	96.4	65.1
Cedar apple rust	Laboratory	45.9	95.7	62.8
Alternaria leaf spot	Laboratory	46.1	95.8	63.5
	Field	37.6	89.4	54.1
Rust	Laboratory	45.7	95.6	62.9
	Field	38.9	90.7	55.6

4.4. Comparison with Existing Methods

To verify the effectiveness of the proposed method, we compared it with several baseline methods (Table 4), including the following: CAM Baseline: directly generating lesion regions using traditional CAM; CCAM: a weakly supervised localization method based on contrastive learning; Grad-CAM: generating CAM by leveraging the class-specific weights of diseased classes in the output layer and the feature maps before the global average pooling layer. Our method first applies probabilistic CAM and then performs feature decoupling.

Figure 5 presents qualitative visualizations of the intermediate localization outputs and final segmentation results produced by the compared methods. As shown in Figure 5, ResNet-CAM is able to identify image regions containing visible disease spots, indicating

that the model learns disease-related visual cues. However, its localization maps remain relatively coarse. Several background pixels, particularly shadow regions, are incorrectly activated, while some diseased regions are not covered. These false-positive and false-negative regions reduce the spatial overlap between the predicted mask and the ground-truth mask, resulting in a relatively low IoU.

In contrast, the proposed framework produces more spatially concentrated lesion localization after probabilistic CAM-guided feature disentanglement. The subsequent pseudo-mask refinement and RL-based prompt optimization further improve the spatial completeness and boundary quality of the predicted lesion regions. These qualitative observations are consistent with the quantitative results reported in Table 4.




Figure 5. Visualization comparison.

Table 4. Comparison of different methods on segmentation performance.

Method	Supervision	mIoU (%)	Precision (%)	Recall (%)
U-net	Full	75.0	84.5	85.4
DeepLab	Full	82.9	89.7	90.5
WSLSS-400	Semi	43.4	74.7	58.5
ResNet-CAM	Weak	19.8	30.4	50.3
Grad-CAM + ExG	Weak	18.0	28.0	48.9
DeepLab + Pseudo	Weak	19.7	30.8	49.7
DeepLab + Pseudo + ExG	Weak	35.7	71.3	51.1
WSLSS-200	Weak	43.1	74.9	57.9
Ours (no RLPO)	Weak	44.1	94.6	55.8
Ours	Weak	45.6	95.4	63.6

4.5. Ablation Study

To analyze the contribution of each component, we conduct ablation experiments (Table 5).

The results show that each component contributes to performance improvement. In particular, the RL-based prompt optimization significantly improves boundary refinement, while PCFD and DLC enhance pseudo-mask quality.

Table 5. Ablation study of probabilistic CAM and feature disentanglement.

PCFD	DLC	RLPO	mIoU (%)	F1-Score (%)
✓	×	×	41.6 ± 0.7	52.9 ± 0.8
×	✓	×	41.2 ± 0.8	52.8 ± 0.9
✓	✓	×	44.1 ± 0.6	55.8 ± 0.7
✓	✓	✓	45.6 ± 0.5	57.1 ± 0.6

4.6. Effect of Prompt Optimization

To validate the effectiveness of the proposed prompt optimization strategy, we compare it with heuristic methods in Table 6 and also compare the impact of the number of prompts on the metrics in Table 7.

Table 6. Comparison of RL prompt strategies.

Prompt Strategy	Description	mIoU (%)	F1-Score (%)
Baseline	Random sampling	42.6 ± 1.0	54.2 ± 1.1
Heuristic + random	Distance-based selection	44.4 ± 0.7	56.1 ± 0.8
RL (Ours)	Learned policy	45.6 ± 0.5	57.1 ± 0.6

Table 7. Effect of number of prompts.

Number of Prompts	mIoU (%)	F1-Score (%)
1	43.1 ± 0.8	53.8 ± 0.9
2	43.9 ± 0.6	54.3 ± 0.7
3	45.6 ± 0.5	57.1 ± 0.6
5	40.5 ± 0.8	52.0 ± 0.7
10	37.9 ± 1.0	49.2 ± 1.1

The proposed RLPO strategy achieves higher mean segmentation performance than the random and heuristic prompt-selection strategies across the repeated runs.

4.7. Visualization Results

Figure 6 shows qualitative segmentation results.




Figure 6. Model effect visualization.

The proposed method produces more complete lesion regions and cleaner boundaries compared with baseline methods. In particular, it effectively suppresses background noise and captures fine-grained lesion structures.

4.8. Cross-Crop Generalization

To further evaluate the generalization ability of the proposed method, we conducted experiments on grape and strawberry leaf disease datasets. The proposed method achieved an mIoU of 37.9% in the cross-dataset evaluation, as reported in Table 8.

Table 8. Cross-crop generalization from apple to grape leaf disease images.

Method	mIoU (%)	Precision (%)	Recall (%)
Ours	37.9 ± 0.8	82.4 ± 1.2	66.3 ± 1.0

Compared with the in-domain performance on apple leaves (45.6%), the performance degradation indicates that the proposed method still suffers from domain shifts across different crops. Differences in leaf appearance, disease characteristics, background, and imaging conditions may affect the quality of disease localization and subsequent SAM prompt optimization. Therefore, the cross-dataset result demonstrates partial transferability rather than strong cross-crop generalization. While the proposed framework can still produce meaningful segmentation on an unseen crop, its performance decreases substantially without target-domain adaptation. Improving crop-invariant representation learning and domain adaptation will be considered in future work.

5. Conclusions

This paper presents a weakly supervised framework for leaf disease segmentation that integrates probabilistic feature disentanglement and reinforcement learning-based prompt optimization. The proposed method addresses key challenges in agricultural image analysis, including incomplete lesion localization, background interference, and the lack of effective prompt selection strategies.

By introducing probabilistic CAM-guided feature disentanglement and a dual-loss constraint, the framework improves the quality and stability of pseudo masks under weak supervision. In addition, the proposed prompt optimization strategy enables adaptive and

automatic selection of informative prompts, allowing a foundation segmentation model to refine lesion boundaries without human intervention.

Experimental results demonstrate that the proposed method achieves competitive performance compared with existing approaches while significantly reducing annotation requirements. Cross-crop experiments further demonstrate that the proposed framework retains meaningful segmentation capability on unseen crops, although a performance degradation is observed under cross-domain conditions. These results indicate partial transferability while highlighting the need for further domain adaptation to improve robustness across different imaging conditions and crop types.

Future work will focus on improving computational efficiency and extending the proposed framework to more diverse agricultural datasets and real-world field conditions.

Author Contributions: Conceptualization, J.Z. and J.J.; methodology, J.J.; software, J.J.; validation, J.J., Y.C., and X.Z.; formal analysis, J.J., W.H., and S.H.; investigation, Y.C., S.H., and W.H.; data curation, J.J., Y.C., and X.Z.; writing—original draft preparation, J.J.; writing—review and editing, J.Z. and J.J.; visualization, J.J.; supervision, J.Z.; project administration, J.Z. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Jiangxi Provincial Natural Science Foundation, grant number 20252BAC240198.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data presented in this study are available on request from the corresponding author.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

SAM	Segment Anything Model
CAM	Class Activation Mapping
RLPO	Reinforcement Learning-based Prompt Optimization
PCFD	Probabilistic CAM-guided Feature Disentanglement
DLC	Dual-Loss Constraint
mIoU	Mean Intersection over Union
CNN	Convolutional Neural Network
MDP	Markov Decision Process

References

1. Toda, Y.; Okura, F. How convolutional neural networks diagnose plant disease. *Plant Phenomics* **2019**, *2019*, 9237136.
2. Zhu, X.; Shao, S.; Shao, S. Prediction of collapsibility of loess site based on artificial intelligence: Comparison of different algorithms. *Environ. Earth Sci.* **2024**, *83*, 101.
3. Zeng, W.; He, M. Rice disease segmentation method based on CBAM-CARAFE-DeepLabv3+. *Crop Prot.* **2024**, *180*, 106665.
4. Luo, Z.; Yang, W.; Yuan, Y.; Gou, R.; Li, X. Semantic segmentation of agricultural images: A survey. *Inf. Process. Agric.* **2024**, *11*, 172–186.
5. Zhang, X.; Cen, C.; Li, F.; Liu, M.; Mu, W. CRFormer: Cross-resolution transformer for segmentation of grape leaf diseases with context mining. *Expert Syst. Appl.* **2023**, *229*, 120324.
6. Chang, Z.; Li, H.; Chen, D.; Liu, Y.; Zou, C.; Chen, J.; Han, W.; Liu, S.; Zhang, N. Crop type identification using high-resolution remote sensing images based on an improved DeepLabv3+ network. *Remote Sens.* **2023**, *15*, 5088.
7. Yu, F.; Yu, X. LEU-Net: A lightweight edge-aware U-Net for medical image segmentation. *IEEE Access* **2026**, *14*, 16334–16346.

8. Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Kaiser, L.; Polosukhin, I. Attention is all you need. In Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS'17), Long Beach, CA, USA, 4–9 December 2017.
9. Liu, X.; Gao, P.; Yu, T.; Wang, F.; Yuan, R.Y. CSWin-UNet: Transformer UNet with cross-shaped windows for medical image segmentation. *Inf. Fusion* **2025**, *113*, 102634.
10. Chan, L.; Hosseini, M.S.; Plataniotis, K.N. A comprehensive analysis of weakly-supervised semantic segmentation in different image domains. *Int. J. Comput. Vis.* **2021**, *129*, 361–384.
11. Liu, X.; Wang, R.; Lai, Y.; Shi, G.; Shao, F.; Hao, F.; Zhang, J.; Shen, J.; Wu, Y.; Zheng, W. Plug-and-play PPO: An adaptive point prompt optimizer making SAM greater. In Proceedings of the 2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Nashville, TN, USA, 11–15 June 2025.
12. Wang, R.; Yang, Z.; Song, Y. OSAM-Fundus: A training-free, one-shot segmentation framework for optic disc and cup in fundus images. *Biomed. Signal Process. Control* **2024**, *100*, 107069.
13. Jung, H.; Oh, Y. Towards better explanations of class activation mapping. In Proceedings of the 18th IEEE/CVF International Conference on Computer Vision (ICCV 2021), Montreal, QC, Canada, 10–17 October 2021.
14. Zhang, W.; Zhang, H.; Zhang, X.; Shen, X.; Zou, L. IP-CAM: Class activation mapping based on importance weights and principal-component weights for better and simpler visual explanations. *Comput. Vis. Image Underst.* **2025**, *261*, 104523.
15. Lin, K.; Gong, L.; Huang, Y.; Liu, C.; Pan, J. Deep learning-based segmentation and quantification of cucumber powdery mildew using convolutional neural network. *Front. Plant Sci.* **2019**, *10*, 155.
16. Tian, L.; Yuan, Y.; En, Q.; Ma, W.; Zhang, G.; Liang, F. Attention-based unsupervised prompt learning for SAM in leaf disease segmentation. *Knowl.-Based Syst.* **2025**, *320*, 113652.
17. Deng, Y.; Xi, H.; Zhou, G.; Chen, A.; Wang, Y.; Li, L.; Hu, Y. An effective image-based tomato leaf disease segmentation method using MC-UNet. *Plant Phenomics* **2023**, *5*, 49.
18. Zhang, Y.; Lu, H.; Ma, G.; Zhao, H.; Xie, D.; Geng, S.; Tian, W.; Sian, K.T.C.L.K. MU-Net: Embedding MixFormer into UNet to extract water bodies from remote sensing images. *Remote Sens.* **2023**, *15*, 3559.
19. Awedat, K.; Alajmi, M. Leveraging U-Net and ASPP for effective fault detection in photovoltaic modules. *Sci. Rep.* **2025**, *15*, 21788.
20. Wei, T.; Chen, Z.; Yang, X.; Chen, S.; Zhang, M.; Huang, Z. A large-scale in-the-wild dataset for plant disease segmentation. *Sci. Data* **2026**, *13*, 205.
21. Guo, Z.; Bian, L.; Wei, H.; Li, J.; Ni, H.; Huang, X. DSNet: A novel way to use atrous convolutions in semantic segmentation. *IEEE Trans. Circuits Syst. Video Technol.* **2025**, *35*, 3679–3692.
22. Wang, K.; Hatzinakos, D. SEFormer: Dual-path conformer neural network is a good speech denoiser. In Proceedings of the 2023 Asia Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC), Taipei, Taiwan, 31 October–3 November 2023.
23. Shafiq, M.; Gu, Z. Deep residual learning for image recognition: A survey. *Appl. Sci.* **2022**, *12*, 8972.
24. Zhang, X.; Li, F.; Zheng, H.; Mu, W. UPFormer: U-shaped perception lightweight transformer for segmentation of field grape leaf diseases. *Expert Syst. Appl.* **2024**, *249*, 123546.
25. Selvaraju, R.R.; Cogswell, M.; Das, A.; Vedantam, R.; Parikh, D.; Batra, D. Grad-CAM: Visual explanations from deep networks via gradient-based localization. *Int. J. Comput. Vis.* **2020**, *128*, 336–359.
26. Dai, W.; Zhu, W.; Zhou, G.; Liu, G.; Xu, J.; Zhou, H.; Hu, Y.; Liu, Z.; Li, J.; Li, L. AISOA-SSFormer: An effective image segmentation method for rice leaf disease based on the transformer architecture. *Plant Phenomics* **2024**, *6*, 0218.
27. Gu, A.; Dao, T. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv* **2023**, arXiv:2312.00752.
28. Zhou, L.; Xiao, Q.; Taha, M.F.; Xu, C.; Zhang, C. Phenotypic analysis of diseased plant leaves using supervised and weakly supervised deep learning. *Plant Phenomics* **2023**, *5*, 22.
29. Xie, J.; Xiang, J.; Chen, J.; Hou, X.; Zhao, X.; Shen, L. C2AM: Contrastive learning of class-agnostic activation map for weakly supervised object localization and semantic segmentation. In Proceedings of the 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), New Orleans, LA, USA, 18–24 June 2022.
30. Xu, Y.; Tang, J.; Men, A.; Chen, Q. EviPrompt: A training-free evidential prompt generation method for adapting segment anything model in medical images. *IEEE Trans. Image Process.* **2024**, *33*, 6204–6215.
31. Zhou, H.; Qiao, B.; Yang, L.; Lai, J.; Xie, X. Texture-guided saliency distilling for unsupervised salient object detection. In Proceedings of the 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Vancouver, BC, Canada, 17–24 June 2023.
32. Siméoni, O.; Sekkat, C.; Puy, G.; Vobecký, A.; Zablocki, É.; Pérez, P. Unsupervised object localization: Observing the background to discover objects. In Proceedings of the 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Vancouver, BC, Canada, 17–24 June 2023.
33. Chen, Y.; Chen, B.; Xian, W.; Wang, J.; Gong, X.; Han, J.; Tao, X. SAM-IAD: Injecting specific knowledge into SAM for industrial anomaly detection. *Knowl.-Based Syst.* **2026**, *338*, 115515.

34. Ma, J.; He, Y.; Li, F.; Han, L.; You, C.; Wang, B. Segment anything in medical images. *Nat. Commun.* **2024**, *15*, 654.
35. Sajitha, P.; Andrushia, A.D.; Anand, N.; Naser, M. A review on machine learning and deep learning image-based plant disease classification for industrial farming systems. *J. Ind. Inf. Integr.* **2024**, *38*, 100572.
36. Osco, L.P.; Wu, Q.; de Lemos, E.L.; Gonçalves, W.N.; Ramos, A.P.M.; Li, J.; Junior, J.M. The segment anything model (SAM) for remote sensing applications: From zero to one shot. *Int. J. Appl. Earth Obs. Geoinf.* **2023**, *124*, 103540.
37. Liu, Y.; Zhu, M.; Li, H.; Chen, H.; Wang, X.; Shen, C. Matcher: Segment anything with one shot using all-purpose feature matching. In Proceedings of the 12th International Conference on Learning Representations (ICLR 2024), Vienna, Austria, 7–11 May 2024.
38. Liu, X.; Shi, G.; Wang, R.; Lai, Y.; Zhang, J.; Sun, L.; Yang, Q.; Wu, Y.; Li, M.; Han, W.; et al. Feature-prompting GBMSeg: One-shot reference guided training-free prompt engineering for glomerular basement membrane segmentation. In Proceedings of the Medical Image Computing and Computer Assisted Intervention (MICCAI 2024), Marrakesh, Morocco, 6–10 October 2024.
39. Kumari, V.; Bansal, V. A systematic review of plant leaf disease detection and classification using machine learning and deep learning techniques. *Int. J. Data Sci. Anal.* **2026**, *22*, 22.
40. Lv, Y.; Zhang, J.; Dai, Y. Weakly-supervised contrastive learning for unsupervised object discovery. *IEEE Trans. Image Process.* **2024**, *33*, 2689–2702.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.