

Article

Detection of Citrus Huanglongbing in Natural Field Conditions Using an Enhanced YOLO11 Framework

Liang Cao ¹, Wei Xiao ¹, Zeng Hu ¹ , Xiangli Li ^{1,*} and Zhongzhen Wu ^{2,*}¹ College of Artificial Intelligence, Zhongkai University of Agriculture and Engineering, Guangzhou 510225, China; cll@zhku.edu.cn (L.C.); xiaowei@zhku.edu.cn (W.X.); huzeng@zhku.edu.cn (Z.H.)² College of Agriculture & Biology, Zhongkai University of Agriculture and Engineering, Guangzhou 510225, China

* Correspondence: lixiangli@zhku.edu.cn (X.L.); wuzhongzhen@zhku.edu.cn (Z.W.)

Abstract

Citrus Huanglongbing (HLB) is one of the most devastating diseases in the global citrus industry, but its early detection under complex field conditions remains a major challenge. Existing methods often suffer from insufficient dataset diversity and poor generalization, and struggle to accurately detect subtle early-stage lesions and multiple HLB symptoms in natural backgrounds. To address these issues, we propose an enhanced YOLO11-based framework, DCH-YOLO11. We constructed a multi-symptom HLB leaf dataset (MS-HLBD) containing 9219 annotated images across five classes: Healthy (1862), HLB blotchy mottling (2040), HLB Zinc deficiency (1988), HLB yellowing (1768), and Canker (1561), collected under diverse field conditions. To improve detection performance, the DCH-YOLO11 framework incorporates three novel modules: the C3k2 Dynamic Feature Fusion (C3k2_DFF) module, which enhances early and subtle lesion detection through dynamic feature fusion; the C2PSA Context Anchor Attention (C2PSA_CAA) module, which leverages context anchor attention to strengthen feature extraction in complex vein regions; and the High-efficiency Dynamic Feature Pyramid Network (HDFPN) module, which optimizes multi-scale feature interaction to boost detection accuracy across different object sizes. On the MS-HLBD dataset, DCH-YOLO11 achieved a precision of 91.6%, recall of 87.1%, F1-score of 89.3, and mAP50 of 93.1%, surpassing Faster R-CNN, SSD, RT-DETR, YOLOv7-tiny, YOLOv8n, YOLOv9-tiny, YOLOv10n, YOLO11n, and YOLOv12n by 13.6%, 8.8%, 5.3%, 3.2%, 2.0%, 1.6%, 2.6%, 1.8%, and 1.6% in mAP50, respectively. On a publicly available citrus HLB dataset, DCH-YOLO11 achieved a precision of 82.7%, recall of 81.8%, F1-score of 82.2, and mAP50 of 89.4%, with mAP50 improvements of 8.9%, 4.0%, 3.8%, 3.2%, 4.7%, 3.2%, and 3.4% over RT-DETR, YOLOv7-tiny, YOLOv8n, YOLOv9-tiny, YOLOv10n, YOLO11n, and YOLOv12n, respectively. These results demonstrate that DCH-YOLO11 achieves both state-of-the-art accuracy and excellent generalization, highlighting its strong potential for robust and practical citrus HLB detection in real-world applications.

Keywords: YOLO; natural environment; HLB; feature fusion; attention mechanism**MSC:** 68T07

Academic Editor: Jonathan Blackledge

Received: 4 June 2025

Revised: 6 July 2025

Accepted: 7 July 2025

Published: 8 July 2025

Citation: Cao, L.; Xiao, W.; Hu, Z.; Li, X.; Wu, Z. Detection of Citrus Huanglongbing in Natural Field Conditions Using an Enhanced YOLO11 Framework. *Mathematics* **2025**, *13*, 2223. <https://doi.org/10.3390/math13142223>

Copyright: © 2025 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) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Citrus Huanglongbing (HLB), also known as citrus greening disease, is one of the most severe threats facing the global citrus industry, particularly in major production regions such as Asia, Africa, and South America. Characterized by rapid transmission

and the difficulty of eradication, HLB has caused significant economic losses worldwide [1,2]. The pathogen is a phloem-restricted, obligate bacterium—*Candidatus Liberibacter* spp.—primarily transmitted by citrus psyllids. Infected trees exhibit symptoms including stunted growth, leaf chlorosis, fruit quality degradation, and, in severe cases, death, thereby posing a major risk to the sustainability and profitability of citrus cultivation [3,4]. Due to the inconspicuous nature of early-stage symptoms, timely and accurate diagnosis is crucial for effective HLB management and control [5].

Currently, mainstream HLB detection methods include field visual inspection, polymerase chain reaction (PCR), spectroscopic techniques, and computer-vision-based automated detection systems [6]. Field visual inspection relies on experienced agronomists or plant protection specialists to identify phenotypic symptoms on leaves, stems, and fruits. While this approach is simple and enables rapid preliminary screening in orchards, it is highly subjective and depends heavily on the operator's expertise, resulting in low accuracy for early or latent infections [7]. PCR technology, by directly detecting the pathogen's nucleic acid, is regarded as one of the most precise diagnostic tools for HLB [8]. It offers high sensitivity and specificity and can identify infections even before symptoms become apparent. However, PCR requires strict laboratory conditions, sophisticated equipment, and high-purity reagents. Its complex procedures and high costs hinder large-scale, in-field deployment and immediate screening. Spectroscopic analysis achieves non-destructive evaluation of plant health by analyzing leaf spectral characteristics [9]. Although this method is adaptable to various environmental conditions and provides high-throughput information, it is susceptible to interference from light variation and background noise. Moreover, it involves complex data processing and feature extraction, necessitating specialized equipment and algorithms. Automated detection based on computer vision has recently attracted extensive attention. Leveraging deep learning and artificial intelligence, these methods enable rapid and efficient identification of symptomatic leaves under complex field conditions [10]. While they have reduced the labor intensity and improved detection efficiency and consistency, challenges remain in terms of insufficient dataset diversity and limited model generalization. In summary, although current detection techniques have advanced the efficiency of HLB diagnosis, they still present challenges such as labor-intensive procedures, high dependence on operator expertise, high costs, and limited adaptability to varying field conditions. These limitations hinder their large-scale application for high-throughput, low-cost, and rapid screening. Therefore, there is an urgent need to develop a novel HLB detection technology that combines efficiency, cost-effectiveness, and robust adaptability to real-world environments.

In recent years, significant progress has been made in the field of citrus HLB detection [11]. Liu et al. [12] proposed a method for leaf classification based on image texture features, which is straightforward and capable of identifying local lesions to some extent. However, this approach relies heavily on manual feature extraction, resulting in poor adaptability to complex backgrounds and early subtle symptoms, and thus exhibits limited generalization capability. Yan et al. [13] enhanced early non-invasive detection by integrating hyperspectral imaging with machine learning. While this method offers high sensitivity, it requires expensive equipment and complex procedures, and its adaptability to field environments is insufficient, making large-scale real-time monitoring difficult. Aswini et al. [14] applied YOLOv7 for automated HLB detection, which significantly improved detection efficiency and automation. Nevertheless, its ability to recognize subtle symptoms and diverse phenotypes in complex environments remains inadequate, and its generalization performance is limited. Li et al. [15] developed YOLOv8-MC, a model primarily aimed at detecting citrus psyllids. While achieving high accuracy and providing technical support for pest monitoring, it cannot directly identify multiple symptomatic fea-

tures of HLB leaves, failing to meet the need for precise disease classification. Lin et al. [16] employed UAV imagery and support vector machines (SVMs) to conduct large-scale HLB monitoring, which is suitable for extensive screening. However, the limited resolution of UAV imagery makes it difficult to detect early-stage subtle lesions, and detection accuracy is greatly affected by environmental factors. Lu et al. [17] improved model robustness in fruit symptom recognition using Mixup data augmentation and convolutional neural networks (CNNs). Despite these advances, the approach focuses on fruit symptoms and demonstrates limited capability for recognizing diverse leaf symptoms under complex field conditions. Table 1 presents a systematic overview of major methods for HLB detection and their corresponding limitations. There has been steady progress in imaging technologies, machine learning approaches, and automation of detection. However, several fundamental challenges still exist. First, many existing methods have limited generalization, often because their datasets are not diverse enough or are overly fitted to controlled conditions. As a result, their performance in real-world field environments is often unstable. Second, these approaches usually depend on manual feature extraction or expensive equipment, making large-scale or low-cost deployment difficult. Third, current systems are often not sensitive enough to detect early or subtle symptoms, which delays timely management. Finally, the complexity of implementation and the lack of adaptability to changing field conditions further restrict their practical application. All these persistent limitations underscore the necessity for developing more robust, efficient, and field-adaptable detection methods, which this study aims to address.

Table 1. Summary of related works on HLB detection and their limitations.

Study	Methodology	Advantages	Limitations
Liu et al. [12]	Image texture classification	Simple; identifies local lesions	Manual feature extraction; weak for complex backgrounds/early symptoms
Yan et al. [13]	Hyperspectral imaging + Machine Learning	High sensitivity; non-destructive	Expensive; complex; not suitable for large-scale field use
Aswini et al. [14]	YOLOv7	Fast; automated	Weak for subtle/diverse symptoms; limited generalization
Li et al. [15]	YOLOv8-MC	Accurate for pest detection	Cannot detect multiple HLB symptoms; insufficient for disease classification
Lin et al. [16]	UAV imaging + SVM	Large-scale monitoring	Low resolution; environment-dependent; poor for early symptoms
Lu et al. [17]	Mixup + CNN	Robust for fruit symptoms	Focused on fruit; limited for leaf symptoms

Meanwhile, deep learning techniques—particularly CNNs—have significantly advanced the development of intelligent detection systems for fruit and vegetable diseases and pests [18,19]. Numerous studies have demonstrated that methods based on object detection frameworks such as YOLO exhibit outstanding advantages in handling the complexities of agricultural field environments. Jin et al. [20] developed a lightweight CO-YOLO model for *Camellia oleifera* fruit detection, which substantially reduced the model's parameter count while maintaining a high detection accuracy of 90.6%. This result highlights the strong lightweight design and suitability of the YOLO architecture for real-time applications in resource-constrained scenarios. Zhu et al. [21] applied an improved YOLOv7 to a mushroom harvesting system, increasing the overall picking success rate to 80%. This indicates that YOLO-based detection technology not only improves the efficiency of agricultural automation but also enhances target detection reliability under complex environmental conditions. Goyal et al. [22] utilized the YOLOv8 model to achieve efficient and automatic identification of weeds in potato fields. The method maintained

a detection accuracy of 100% across various weed species and under different natural lighting conditions, demonstrating the robustness and high precision of the YOLO family of models in practical field scenarios, which made them highly suitable for large-scale automated farmland management. Shi et al. [23] proposed the YOLO11-CGB approach, which integrates feature enhancement strategies with the YOLO framework, significantly improving the detection of Chinese cabbage seedlings in challenging field conditions, with an average precision of 97.0%. This work shows that customized network structures targeting agricultural object characteristics can further enhance the recognition ability of YOLO models, especially for low-contrast and densely packed targets. Niu et al. [24] improved a tea bud detection algorithm based on YOLO11, increasing the mAP50 by 1.4% and effectively enhancing the model's capability to perceive small objects. This demonstrates the potential of the YOLO architecture for optimizing small object detection and provides a technical reference for improving the detection of early-stage lesions and other fine-grained targets. Li et al. [25] introduced the YOLO11-GS model, incorporating an angle regression optimization strategy for precise grape fruit detection (achieving an accuracy of 87.7% and mAP50 of 89.5%). This model effectively addressed the challenge of detecting fruits at varying growth stages and spatial orientations, offering a robust solution for complex field scenarios involving multi-angle and deformed targets. Collectively, these examples illustrate that YOLO-based models, owing to their end-to-end architecture, fast detection speed, and strong adaptability to diverse environments and small objects, have become the mainstream technical approach for intelligent detection of agricultural diseases and pests. Their superior environmental robustness, small object recognition capability, model efficiency, and deployment flexibility provide a solid foundation for large-scale field disease monitoring and automated control. Moreover, these advancements offer important theoretical and methodological insights for further optimizing the YOLO framework to improve the detection of multiple HLB phenotypes, early symptoms, and environmental variability, as pursued in the present study.

A comprehensive analysis of existing research reveals that although deep learning approaches have achieved remarkable progress in the detection of fruit and vegetable diseases, two key challenges persist in the practical detection of citrus HLB: First, current datasets lack sufficient diversity in terms of geography, seasonality, imaging angles, and environmental conditions. This limitation prevents models from fully learning the symptom characteristics of HLB across multiple phenotypes and scenarios, thereby constraining their generalization capability. Second, most existing models exhibit inadequate feature extraction and differentiation abilities for subtle early-stage lesions and various symptomatic types under complex natural backgrounds. This shortcoming hampers the practicality of early, accurate in-field identification and automated multi-symptom diagnosis of HLB.

To address these challenges, this study proposes a novel HLB recognition method based on an enhanced YOLO11 architecture, specifically designed for natural field conditions. The goal is to further improve detection accuracy and robustness in complex environments. This study specifically targets the automated detection of multiple symptomatic features of citrus HLB on leaves under diverse and realistic field conditions, using deep-learning-based computer vision methods. Compared with existing HLB detection approaches, our method offers clear advantages in three main aspects. First, the enhanced YOLO11 framework achieves more robust generalization, maintaining high detection accuracy across diverse and challenging field conditions. Second, it significantly improves the identification of early-stage and multiple symptomatic features, even under complex backgrounds or subtle symptom expression. Third, the proposed model demonstrates greater adaptability and practical value for real-world agricultural deployment. These strengths are validated through comprehensive comparative experiments, where our method con-

sistently surpasses state-of-the-art models in detection precision, robustness, and field applicability. The main contributions of this work are as follows:

- We constructed the Multi-Symptom HLB Leaf Dataset (MS-HLBD), which encompasses samples from multiple regions, various shooting angles, different seasons, and diverse time periods. This significantly enhances the diversity of the dataset, overcoming the limitations of existing datasets in terms of regional and seasonal coverage, and provides a richer foundation of image features for model training and evaluation.
- We propose the DCH-YOLO11 model, which comprehensively improves the detection capability for multiple HLB symptom types and enhances overall generalization performance. Building upon the original YOLO11 framework, three key innovative modules are integrated: the C3k2 Dynamic Feature Fusion (C3k2_DFF) module strengthens the interaction between global and local features, thereby improving the detection of subtle early-stage lesions; the C2PSA Context Anchor Attention (C2PSA_CAA) module highlights disease details in complex leaf vein regions, increasing the model's ability to distinguish various symptom types under challenging backgrounds; and the High-efficiency Dynamic Feature Pyramid Network (HDFPN) module optimizes multi-scale feature fusion, enhancing detection accuracy and robustness across different object scales. These targeted enhancements address existing shortcomings in feature extraction, fine-grained recognition, and multi-scenario adaptability.

Nevertheless, it should be noted that the generalization ability of the proposed model under extreme lighting conditions, rare citrus cultivars, or highly atypical symptom presentations remains to be further validated. In addition, the computational adaptability for real-time deployment in resource-constrained scenarios still requires optimization.

2. Materials and Methods

2.1. Data Collection

The dataset used in this study was collected from citrus orchards located in Conghua District, Nansha District, and Zengcheng District of Guangzhou, Guangdong Province, China. The dataset was sampled from different geographical areas to enhance its representativeness and diversity. Based on this collection, the MS-HLBD was constructed, consisting of healthy citrus leaves, citrus leaves infected with HLB (including blotchy mottling, Zinc-deficiency symptoms, and yellowing) and leaves affected by citrus canker.

To obtain high-quality images rich in detail, a handheld high-resolution imaging device was used for close-range image acquisition (5–20 cm distance). Data collection was carried out from December to the following March to cover the phenotypic changes of citrus plants across different growth stages and disease progression phases. To ensure data diversity and representativeness, samples were taken from different trees and canopy positions, as well as leaves oriented in various directions. Moreover, images were captured at different times of day and under various weather conditions, such as morning, noon, evening, sunny, and cloudy settings, to introduce natural variability in lighting, background textures, and leaf postures. During the photography process, both composition and shooting angles were adjusted appropriately to minimize adverse effects such as backlighting, glare, and image blur. Images with poor quality, including severe blur, overexposure, or incomplete leaf regions, were excluded. Ultimately, 4299 high-quality images were retained to form the initial version of the MS-HLBD. Figure 1 presents examples of typical diseased leaf images collected in the dataset. MS-HLBD exhibits substantial diversity and representativeness in terms of disease types, environmental conditions, and growth stages, providing a solid data foundation for subsequent model training, validation, and testing.




Figure 1. Sample images from the Multi-Symptom HLB Leaf Dataset.

In this study, we did not simply merge all leaf symptoms of HLB into a single class. Instead, we subdivided them into three subclasses based on observable phenotypes in the field: HLB blotchy mottling, HLB Zinc deficiency, and HLB yellowing. This subdivision strategy holds clear practical significance. From the perspective of agricultural management, differentiating symptom types can directly support field decision-making. Blotchy mottling is typically the earliest leaf symptom of HLB and requires immediate enhanced monitoring and marking of infected trees [26]. Zinc-deficiency symptoms indicate the need for mineral supplementation or targeted fertilization, which can temporarily improve tree vigor and delay yield decline [27]. Yellowing symptoms generally signal that the plant has entered a declining stage; current global best practices recommend early removal of symptomatic trees and regional vector (psyllid) control to minimize subsequent inoculum pressure in the orchard [28]. Furthermore, from a computer vision perspective, these three symptom types differ significantly in terms of leaf vein texture, lesion morphology, and color gradients. Explicit subclassification enables the model to learn richer and more discriminative features and also enhances its ability to distinguish HLB from other citrus diseases. In summary, the symptom subdivision strategy adopted in this study not only aligns with agronomic management needs but also facilitates high-quality feature learning of subtle lesions by the model, thereby increasing its practical applicability.

2.2. Data Augmentation and Dataset Preparation

To further improve the accuracy of disease classification and ensure the reliability of annotations, real-time fluorescence quantitative PCR (qPCR) was performed to verify pathogen types [29]. Based on molecular testing results, citrus leaves were categorized by disease type, providing the basis for subsequent image augmentation, model training, and evaluation. From the raw images, 4299 images were selected for inclusion in MS-HLBD based on the quality and availability of annotatable regions. To enhance the model's robustness and accuracy, various data augmentation strategies were applied, including

horizontal flipping, vertical flipping, random Gaussian blur, cropping, and noise addition. Examples of data augmentation effects are shown in Figure 2. After augmentation, the dataset size was expanded to 9219 images. Table 2 summarizes the distribution of images before and after data augmentation.




Figure 2. Sample images showing data augmentation effects.

Table 2. Image distribution before and after data augmentation.

Index	Classes	Abbreviation	No Data Augmentation	Data Augmentation
1	Healthy	HL	868	1862
2	HLB blotchy mottling	HLB_bm	954	2040
3	HLB Zinc deficiency	HLB_Zd	924	1988
4	HLB yellowing	HLB_y	824	1768
5	Canker	Canker	729	1561
Total			4299	9219

3. Experimental Methods

3.1. Improved YOLO11 Model

In this study, after extensive evaluation from multiple perspectives and comparative experiments with various baseline models, YOLO11 was ultimately selected as the base model. YOLO11 inherits the advantages of the YOLO series, such as high-speed and high-precision object detection, and incorporates advanced neural network architectures and optimization strategies to support robust multi-scale feature representation [30–32].

However, when applied to the detection of HLB-infected citrus leaves in practical scenarios, YOLO11 still exhibits several limitations. For instance, the model is easily disturbed by cluttered backgrounds in complex field environments, which often results in missed detections or false positives. Moreover, its feature extraction and fusion capabilities for early subtle lesions are insufficient, resulting in limited accuracy for early symptom recognition. In addition, given the diverse types of citrus diseases and the similarity of symptoms, YOLO11's fine-grained discrimination capability remains relatively weak, making it difficult to reliably distinguish HLB from other leaf diseases.

To address these issues, an improved version of YOLO11, referred to as DCH-YOLO11, is proposed in this study. First, the C3k2_DFF module was introduced to replace the original C3k2 module in the backbone, which enhances the fusion of global and local

features and compensates for the model's deficiency in capturing early subtle disease characteristics. Second, the C2PSA_CAA module was incorporated to improve the spatial feature representation. This module integrates a context anchor attention mechanism to highlight fine-grained details in complex vein regions, thereby significantly boosting the model's ability to discriminate subtle disease features against cluttered backgrounds. Finally, the HDFPN module was constructed to optimize the original multi-scale feature fusion strategy of YOLO11, which enables more effective integration of high-level semantic features with low-level localization features, thereby enhancing detection performance across different scales and reducing interference from environmental background noise. The overall architecture of the DCH-YOLO11 model is illustrated in Figure 3.




Figure 3. Architecture of the proposed DCH-YOLO11 model.

3.1.1.1. C3k2_DFF

The original C3k2 module in YOLO11 employs simple residual addition for feature fusion. However, this approach lacks the ability to adaptively distinguish the importance of different channels or spatial locations, which may lead to important information being overlooked or redundant information being retained during feature fusion. This limitation constrains the model's adaptability to complex scenarios.

To address this issue, this study incorporates the Dynamic Feature Fusion (DFF) mechanism proposed by Yang et al. [33] and redesigns the original C3k2 module into a new version named C3k2_DFF. The C3k2_DFF module adaptively fuses multi-scale features using global channel and spatial attention mechanisms, thereby enhancing its ability to capture subtle lesion features. Leveraging the DFF mechanism significantly improves the model's capacity to extract and fuse early subtle disease features. The architecture of the improved C3k2_DFF module is illustrated in Figure 4, which supports two operating modes ($C3k = \text{False}$ and $C3k = \text{True}$) to meet varying application needs and network configurations.




Figure 4. Structure of the C3k2_DFF module.

The DFF module is designed to adaptively fuse internal features across different scales by leveraging global contextual information. Specifically, given two feature maps $F_1^l \in \mathbb{R}^{C \times H \times W \times D}$, and $F_2^l \in \mathbb{R}^{C \times H \times W \times D}$ (C denotes the number of channels; H , W and D denote height, width, and depth), they are concatenated along the channel dimension to form $F^l \in \mathbb{R}^{2C \times H \times W \times D}$. This concatenation operation can be expressed as

$$F^l = \text{Concat}\left([F_1^l; F_2^l]\right) \quad (1)$$

To ensure that the module can continuously and effectively utilize the fused features, a channel compression mechanism is needed to reduce the number of channels in the concatenated feature map F^l back to the original number C . Unlike traditional channel compression methods that rely on simple convolution operations, the DFF module introduces a global channel descriptor w_{ch} to dynamically retain important features and suppress redundant information. Specifically, the extraction process of the global channel descriptor w_{ch} is as follows: First, global average pooling (AvgPool) is applied to F^l to compute channel-wise statistics; then, a 1×1 convolutional layer extracts intermediate representations; finally, a Sigmoid activation normalizes the output into an attention weight vector. This attention weight vector w_{ch} is then multiplied element-wise with F^l , followed by a $1 \times 1 \times 1$ convolution to select more critical features, thereby completing the compression along the channel dimension. The detailed process can be formulated as

$$w_{ch} = \text{Sigmoid}\left(\text{Conv}_1\left(\text{AvgPool}\left(F^l\right)\right)\right) \quad (2)$$

$$F_{ch}^l = \text{Conv}_1\left(w_{ch} \otimes F^l\right) \quad (3)$$

After the channel dimension compression, the DFF module further models the spatial dependencies among features to highlight specific regions in the feature map. This is achieved by first applying a $1 \times 1 \times 1$ convolution separately to F_1^l and F_2^l and then summing the outputs. A Sigmoid activation function is applied to normalize the result to the range

[0, 1], generating a spatial attention weight w_{sp} . Finally, the spatial attention weight w_{sp} is applied element-wise to the channel-calibrated feature map F_{ch}^l , resulting in the final fused feature map $\hat{F}^l$. This process can be formulated as

$$w_{sp} = \text{Sigmoid}\left(\text{Conv}_1\left(F_1^l\right) \oplus \text{Conv}_1\left(F_2^l\right)\right) \quad (4)$$

$$\hat{F}^l = w_{sp} \otimes F_{ch}^l \quad (5)$$

3.1.2. C2PSA_CAA

Although the original C2PSA module in YOLO11 effectively enhances feature representation along the channel dimension, it still exhibits limitations in capturing disease-related fine details.

To address this issue, this study introduces the Context Anchor Attention (CAA) mechanism proposed by Cai et al. [34] and constructs a new module, C2PSA_CAA, to further enhance the model's ability to accurately capture fine-grained HLB features in complex field environments. The CAA module is particularly designed to extract elongated texture features, which aligns well with the typical symptoms of HLB-infected citrus leaves, namely, changes in leaf vein textures. Integrating the CAA attention mechanism allows the model to more effectively focus on subtle variations in local diseased leaf veins, thereby improving its sensitivity to early symptoms and overall recognition accuracy. The structure of the C2PSA_CAA module is illustrated in Figure 5.




Figure 5. Structure of the C2PSA_CAA module.

The C2PSA_CAA module first processes the input feature map through a convolutional layer and then partitions the channels into two branches. One branch sequentially passes through multiple PSA blocks, each of which integrates the CAA attention mechanism and a feed-forward network structure. Feature enhancement and fusion are achieved through element-wise residual connections within each PSA block. Finally, the two branches are concatenated along the channel dimension and fused via a convolutional layer to generate the output.

Integrating the C2PSA_CAA module significantly improves the model's ability to capture disease-specific features, providing strong support for improving the subsequent disease recognition performance. The core idea of the CAA module is to replace traditional large 2D convolutional kernels with lightweight horizontal and vertical depthwise convolutions (DWConv), thereby significantly expanding the receptive field in a more efficient manner. Specifically, the CAA module first applies a 7×7 AvgPool operation on the input feature map to capture local contextual information. This is followed by a 1×1 convolution to preliminarily fuse the channel information, formulated as

$$F_{pool}^l = \text{Conv}_{1 \times 1}(\text{AvgPool}(F^l)) \quad (6)$$

Subsequently, 1×11 and 11×1 strip DWConvs are employed sequentially to capture long-range dependencies in the horizontal and vertical directions, respectively. Finally, a 1×1 convolution followed by a Sigmoid activation function is used to generate the attention weights, adaptively enhancing the feature representations in key regions. The detailed process can be expressed as

$$F_h^l = \text{DWConv}_{1 \times 11}(F_{pool}^l) \quad (7)$$

$$F_v^l = \text{DWConv}_{11 \times 1}(F_h^l) \quad (8)$$

$$A^l = \text{Sigmoid}(\text{Conv}_{1 \times 1}(F_v^l)) \quad (9)$$

3.1.3. HDFPN

The original YOLO11 model utilizes a top-down Feature Pyramid Network (FPN) to fuse features across different scales. However, the simple additive fusion strategy fails to fully exploit the complementary advantages of high-level semantics and low-level localization, which may lead to suboptimal multi-scale feature fusion and reduced detection performance for objects of varying sizes.

To address these limitations, this study draws inspiration from the High-Level Screening-feature Fusion Pyramid Network (HSFPN) proposed by Chen et al. [35] and develops an improved structure called the HDFPN. The architecture of the improved HDFPN is illustrated in Figure 6.




Figure 6. Structure of the HDFPN module.

Firstly, this study employs a channel attention (CA) mechanism to adaptively perform channel-wise weighting on high-level feature maps outputted by the backbone network. Specifically, the CA module processes the input feature maps $S_i (i = 3, 4, 5)$ by simultaneously applying global maximum pooling (MaxPool) and global AvgPool. The two pooled features are then element-wise added and activated by a Sigmoid function to generate

channel attention weights α , which are subsequently multiplied with the original input feature maps S_i to realize the adaptive emphasis on important channels. This process can be formulated as

$$\alpha = \text{Sigmoid}(\text{MaxPool}(S_i) + \text{AvgPool}(S_i)), \quad (i = 3, 4, 5) \quad (10)$$

$$P_i = \text{Conv}_{1 \times 1}(\alpha \otimes S_i), \quad (i = 3, 4, 5) \quad (11)$$

Secondly, the Select Feature Fusion (SFF) module initially upsamples the high-level feature map P_5 through Transposed Convolution to obtain $\hat{F}_{\text{high}}$, ensuring spatial resolution matching. After upsampling, the CA module generates a spatial attention weight F_{att} based on $\hat{F}_{\text{high}}$, which is then element-wise multiplied with the lower-level feature maps $P_i (i = 3, 4)$. Subsequently, the fused high-level feature $\hat{F}_{\text{high}}$ and the recalibrated lower-level features are added to obtain the final fused feature map N_i . The detailed process is as follows:

$$\hat{F}_{\text{high}} = \text{TConv}(P_5) \quad (12)$$

$$F_{\text{att}} = \text{CA}(\hat{F}_{\text{high}}) \quad (13)$$

$$N_i = (F_{\text{att}} \otimes P_i) + \hat{F}_{\text{high}}, \quad (i = 3, 4) \quad (14)$$

Finally, to further enhance the network's ability to capture and represent fine-grained disease features after multi-scale feature fusion, a dynamic fine-grained feature extraction module named C3k2_DFF is added to the output feature maps N_i . The C3k2_DFF module integrates the DFF mechanism, adaptively combining features from different spatial and channel dimensions. During dynamic recalibration, feature selection weights are adaptively adjusted, leading to higher-quality disease feature representations. The final feature extraction process can be expressed as

$$\hat{N}_i = \text{C3k2_DFF}(N_i), \quad (i = 3, 4, 5) \quad (15)$$

3.2. Experimental Platform and Model Evaluation Metrics

3.2.1. Experimental Environment and Parameter Settings

To validate the effectiveness of the proposed improvements, the enhanced MS-HLBD dataset was utilized for model training, validation, and testing phases. The dataset was split at a ratio of 8:1:1, resulting in 7380 images for training, 919 images for validation, and 920 images for testing.

The experimental setup was as follows: Windows 11 operating system (64-bit), Intel i7-14700KF CPU, 48 GB RAM, and an NVIDIA RTX 4080 SUPER GPU with 16 GB of VRAM. The model was developed and tested using Python 3.11 and the PyTorch 2.5.1 deep learning framework. According to the hardware configuration, the settings of the pre-training hyperparameters are summarized in Table 3.

Table 3. Hyperparameters and experimental settings used for all model training and evaluation on both the MS-HLBD dataset and the external public HLB dataset.

Parameter	Value
Learning rate	0.01
Momentum	0.937
Weight decay	0.0005
Batch size	16
Epochs	200

3.2.2. Evaluation Metrics

Upon completing model training and inference, a rigorous quantitative evaluation of the model performance is required to assess its effectiveness and robustness in the object detection task. Common evaluation metrics for object detection include Precision (P), Recall (R), F1-score ($F1$), Accuracy (A), Average Precision (AP), and mean Average Precision (mAP). In addition, inference speed is measured by Frames Per Second (FPSs), model size by Parameters, and computational complexity by Floating Point Operations (FLOPs). The calculation formulas for P , R , $F1$, and A are defined as follows:

$$P = \frac{TP}{TP + FP} \quad (16)$$

$$R = \frac{TP}{TP + FN} \quad (17)$$

$$F1 = \frac{2 \times P \times R}{P + R} \quad (18)$$

$$A = \frac{TP + TN}{TP + TN + FP + FN} \quad (19)$$

In the formulas, true positive (TP) refers to positive samples correctly predicted as positive, true negative (TN) refers to negative samples correctly predicted as negative, false positive (FP) refers to negative samples incorrectly predicted as positive, and false negative (FN) refers to positive samples incorrectly predicted as negative. P evaluates the proportion of correctly predicted positive instances among all predicted positive instances. R measures the proportion of correctly detected positive instances among all actual positive instances. $F1$, calculated as the harmonic mean of P and R , is often used as a comprehensive metric to balance the two. A measures the proportion of correctly predicted samples among all samples.

Considering the imbalanced distribution of the classes within the MS-HLBD, relying solely on metrics such as P , R , or A may not adequately reflect the overall classification performance. To address this issue, the Matthews Correlation Coefficient (MCC) is introduced as an additional evaluation metric. MCC comprehensively evaluates classification performance by considering all four components of the confusion matrix (TP , TN , FP , and FN) and is widely acknowledged for its robustness against class imbalance. The MCC is calculated as follows:

$$MCC = \frac{TP \times TN - FP \times FN}{\sqrt{(TP + FP)(TP + FN)(TN + FP)(TN + FN)}} \quad (20)$$

The MCC value ranges from -1 to $+1$, where $+1$ indicates perfect prediction, 0 represents random prediction, and -1 denotes complete disagreement between predictions and observations. The Precision–Recall (PR) curve illustrates the relationship between P and R , with R on the x -axis and P on the y -axis. AP is computed by integrating the P across R values from 0 to 1 , representing the model's overall detection performance under different threshold settings. For multi-class detection tasks, the mAP is obtained by averaging the AP values across all categories, providing an overall measure of the model's detection effectiveness. The formulas for calculating AP and mAP are as follows:

$$AP = \int_0^1 P \cdot (R) dR \quad (21)$$

$$mAP = \frac{1}{N} \sum_{i=1}^N AP_i \quad (22)$$

where N denotes the total number of detection categories, which is set to 5 in this study, and AP_i represents the AP of the i -th class. FPS measures the runtime efficiency of the model in practical applications, indicating the number of image frames that the model can process per second. It serves as a critical metric for evaluating whether the model is suitable for real-time detection scenarios. Parameters indicate the model's complexity and storage demand, whereas FLOPs quantify the computational cost during inference. Lower values of Parameters and FLOPs contribute to faster runtime and improved energy efficiency, making the model more suitable for deployment in resource-constrained environments. All evaluation metrics were measured under the same hardware environment and input dimensions as those used for baseline models, ensuring result comparability.

4. Results and Analysis

4.1. DCH-YOLO11 Model Training Results

Figure 7 presents the loss function changes and the trends of various evaluation metrics during the training process of the proposed improved YOLO11 model (DCH-YOLO11). As shown in Figure 7, both the training and validation losses decrease significantly with the increase in training iterations. The model's P, R, mAP50, and mAP50-95 all reach their optimal values in the later stages of training. The P and R curves remain consistently high with minimal fluctuations, suggesting that the improved DCH-YOLO11 model possesses strong robustness and generalization in detection tasks.




Figure 7. Training and validation loss and evaluation curves of the DCH-YOLO11 model during 200 epochs on the MS-HLBD dataset. Metrics shown include P, R, mAP50, and mAP50-95.

Figure 8a shows the F1 curve of the improved model on the test set, which is used to assess the model's detection balance performance at different confidence thresholds. Overall, the F1 for most categories increases rapidly at lower confidence levels, maintains a high level within the range of 0.3 to 0.9, and then gradually decreases. This trend shows that the model effectively balances P and R within a moderate confidence range. Figure 8b presents the PR curve, which is used to analyze the variation in detection P across different R rates. It can be observed that the PR curves for the five target categories are predominantly concentrated in the upper-right quadrant of the graph, with small fluctuations, reflecting the model's excellent stability and generalization ability. Notably, even at high R levels, the model maintains high P, highlighting its strong capability in suppressing false negatives.




Figure 8. Performance curves of DCH-YOLO11 on the MS-HLBD test set: (a) F1 versus confidence threshold; (b) PR curves for the 5 target classes.

To comprehensively evaluate the improvements in class-wise recognition performance brought by the DCH-YOLO11 model, we conducted a comparative analysis using confusion matrices for both the original YOLO11n and the enhanced DCH-YOLO11 (as shown in Figure 9). The confusion matrices were column-normalized, meaning each column represents the true class, and the values indicate the proportion of true samples predicted as each class. Overall, the DCH-YOLO11 model significantly improved recognition accuracy across all categories, with a particularly notable advantage in hl, which are often confused with background. Specifically, the R rate for the hl category increased from 0.70 with the original YOLO11n to 0.77 with DCH-YOLO11, marking a substantial improvement of 7 percentage points. Simultaneously, the proportion of background samples misclassified as hl decreased from 0.23 to 0.18, a reduction of 5 percentage points. These results demonstrate that the improved model not only enhances the P of feature extraction for healthy leaves but also more effectively reduces background interference in hl detection. Additionally, the recognition accuracy of other disease categories (such as HLB_Zd, HLB_bm, and HLB_y) was also improved to varying degrees, with an average increase of approximately 3 percentage points. These findings further validate the effectiveness of DCH-YOLO11 in improving the capture of subtle features and reducing both false positives and false negatives.




Figure 9. Confusion matrices on the MS-HLBD test set for (a) the original YOLO11n baseline and (b) the proposed DCH-YOLO11 model.

4.2. Comparative Experiments

To further validate the performance and effectiveness of the proposed DCH-YOLO11 model, this study conducted comparative experiments with several leading and high-performance object detection models, using the same dataset and experimental conditions.

These models include the traditional two-stage model Faster-RCNN [36] and the advanced one-stage models SSD [37] and Transformer-based RT-DETR [38], as well as a series of classic YOLO models (YOLOv7-tiny [39], YOLOv8n [40], YOLOv9-tiny [41], YOLOv10n [42], YOLO11n [43], and YOLOv12n [44]). The performance metrics, MCC values, and parameter counts for each model are summarized in Table 4.

The results demonstrate that the proposed DCH-YOLO11 model outperforms all the compared models across several key evaluation metrics, achieving leading overall performance. Specifically, DCH-YOLO11 achieves P, R, and F1 of 91.6%, 87.1%, and 89.3, respectively, and an MCC of 0.741. The high MCC indicates not only strong F1 performance but also a balanced trade-off between true positive and true negative predictions, reflecting robustness against class imbalance. Compared to YOLO11n (MCC 0.684), DCH-YOLO11 shows a notable MCC improvement of 0.057, which aligns with the observed gains in F1, thereby further validating the model's enhanced discriminative capability. Moreover, the mAP50 and mAP50-95 scores of DCH-YOLO11 are 93.1% and 91.5%, respectively, which are the highest among all the compared models, indicating stronger accuracy and generalization ability in detecting objects across different categories. Compared to the original YOLO11n, DCH-YOLO11 achieves improvements of 0.2%, 4.3%, 2.4, 1.8%, 1.6%, and 0.057 in P, R, F1, mAP50, mAP50-95, and MCC, respectively, while maintaining a similar model parameter count (~2.99 M), further optimizing detection performance.

Notably, on the MS-HLBD dataset, RT-DETR achieves the highest P (94.1%) among all compared models. It tends to favor high-confidence predictions and suppress false positives, which is characteristic of a more conservative detection strategy. However, this conservative approach also leads to more missed detections, resulting in lower R (80.1%) and overall balance (MCC 0.509). Such a trade-off between P and R is typical for transformer-based detectors, especially when applied to imbalanced or complex agricultural datasets.

Table 4. Performance comparison of DCH-YOLO11 and other models on the MS-HLBD dataset. The table reports P, R, F1, MCC, mAP50, mAP50-95, and Parameters. Best results are highlighted in bold.

Model	P/%	R/%	F1	MCC	mAP50/%	mAP50-95/%	Parameters
Faster-RCNN	45.7	86.3	58.8	0.180	79.5	64.5	28,316,308
SSD	91.5	78.3	84.4	0.531	84.3	72.7	4,074,532
RT-DETR	94.1	80.1	86.5	0.509	87.8	85.8	31,994,015
YOLOv7-tiny	92.3	80.8	86.2	0.622	89.9	85.6	6,018,420
YOLOv8n	92.8	82.4	87.3	0.663	91.1	90.3	3,006,623
YOLOv9-tiny	93.5	82.5	87.7	0.657	91.5	90.8	2,618,510
YOLOv10n	93.2	81.9	87.2	0.714	90.5	88.8	2,696,366
YOLO11n	91.4	82.8	86.9	0.684	91.3	89.9	2,583,127
YOLOv12n	91.4	83.1	87.1	0.675	91.5	89.8	2,557,703
DCH-YOLO11	91.6	87.1	89.3	0.741	93.1	91.5	2,995,667

The superior performance of DCH-YOLO11 in R, F1, MCC, and mAP can be attributed to the effective capture of early subtle features by the proposed DFF module. Additionally, the CAA mechanism significantly enhances the model's ability to finely detect the leaf vein regions. Meanwhile, the HDFPN module optimizes the feature pyramid structure, improving the model's ability to balance detection accuracy and interference resistance across different scales.

Figure 10 directly presents a comparison of the overall performance of different models. This radar chart compares the performance of 10 models across multiple evaluation metrics, including P, R, F1, MCC, mAP50, mAP50-95, and Parameters. In the radar chart, each curve represents one model, and the closer the intersection of the curve to the edge of the chart, the better the model's performance on the corresponding metric. The larger the area enclosed by the curve, the stronger the overall performance of the model. It is evident that

DCH-YOLO11 outperforms the other models in overall performance, particularly with its advantages in R, F1, MCC, mAP50, and mAP50-95, which are crucial for the model's deployment in real-world applications.




Figure 10. Radar chart comparing 10 detection models on the MS-HLBD dataset under identical conditions. Metrics: P, R, F1, MCC, mAP50, mAP50-95, and Parameters. A larger enclosed area indicates better overall performance across the evaluated metrics.

To comprehensively demonstrate the detection advantages of DCH-YOLO11, this study conducted a comparative visualization analysis of detection results from multiple models. As illustrated in Figure 11, where false positives and missed detections are marked with red arrows, distinct performance differences are observed across models.

In subfigure (a), YOLO11 misclassifies hl as canker. In (b), YOLO11 fails to detect HLB_y, while YOLOv12 incorrectly classifies two hl instances as HLB_bm. In (c), both YOLO11 and YOLOv12 miss the HLB_y lesion. In (d), the same two models again fail to detect hl. In (e), YOLOv8 misses HLB_y, and YOLOv10 not only overlooks both HLB_bm and HLB_y but also incorrectly merges two separate canker instances in the lower-right corner into a single bounding box. While other models can detect most targets, DCH-YOLO11 shows clear superiority in localization accuracy and detection confidence, particularly under complex backgrounds and subtle symptom conditions.

In summary, both the quantitative evaluations and visual comparison results consistently demonstrate that the proposed DCH-YOLO11 model exhibits notable advantages in terms of P, R, and robustness. These findings confirm the effectiveness of the proposed architectural enhancements and highlight the model's potential for practical deployment in complex field environments.



Figure 11. Sample detection results of different models on the MS-HLBD test images. Subfigures (a–e) present five representative test images selected from the MS-HLBD test set, illustrating typical detection performance and error cases for different models. Red arrows indicate missed detections or misclassifications.

4.3. Ablation Experiments

To validate the effectiveness of each improvement module in the DCH-YOLO11 model, a series of ablation experiments were conducted on the MS-HLBD dataset. In these experiments, the original YOLO11n model served as the baseline, with the C3k2_DFF, C2PSA_CAA, and HDFPN modules added incrementally. The performance metrics for the models with the improved modules were compared with the baseline model, as summarized in Table 5.

Specifically, introducing the C3k2_DFF module alone increased R to 83.7% and improved the F1 to 87.3, confirming its effectiveness in fine-grained feature extraction. The C2PSA_CAA module significantly enhanced P (92.5%), highlighting its strong ability to guide the model to focus on contextual structures. Introducing the HDFPN module yielded the highest P of 93.3%, demonstrating its advantage in multi-scale feature fusion. When any two modules were combined, the model's performance improved further, showing clear complementary characteristics. Notably, the combination of C3k2_DFF and C2PSA_CAA led to the most significant improvements in R (84.5%) and mAP50 (92.6%). When all three modules were integrated, the model achieved optimal performance across all metrics, with

P, R, F1, mAP50, and mAP50-95 reaching 91.6%, 87.1%, 89.3, 93.1%, and 91.5%, respectively. These results represent improvements of 0.2%, 4.3%, 2.4%, 1.8%, and 1.6% over the original YOLO11n, respectively. At the same time, the MCC reached its highest value of 0.741, further demonstrating that the proposed modules improve not only accuracy-related metrics but also the overall robustness and reliability of the model's classification. Furthermore, the final model's parameter count was controlled at 2.99M, maintaining a good lightweight characteristic.

Table 5. Ablation study results on the MS-HLBD dataset using YOLO11n as baseline. “✓” indicates that the module is used, while “-” denotes that the module is not used. Best performance is highlighted in bold.

C3k2_DFF	C2PSA_CAA	HDFPN	P/%	R/%	F1	MCC	mAP50/%	mAP50-95/%	Parameters
-	-	-	91.4	82.8	86.9	0.684	91.3	89.9	2,583,127
✓	-	-	91.3	83.7	87.3	0.702	92.0	90.5	2,779,323
-	✓	-	92.5	83.5	87.8	0.699	91.8	90.0	2,569,559
-	-	✓	93.3	83.1	87.9	0.707	92.0	90.6	2,813,039
✓	✓	-	91.3	84.5	87.8	0.692	92.6	90.3	2,765,755
✓	-	✓	91.7	84.5	88.0	0.715	92.4	90.7	3,009,235
-	✓	✓	92.2	83.9	87.9	0.706	92.1	90.4	2,799,471
✓	✓	✓	91.6	87.1	89.3	0.741	93.1	91.5	2,995,667

It is noteworthy that among all ablation settings, introducing HDFPN alone yields the highest P (93.3%), suggesting that enhanced multi-scale feature fusion is especially effective in suppressing false positives and focusing on clear, well-defined lesions. However, the R (83.1%) in this setting remains suboptimal, indicating that some subtle or atypical diseased areas are still missed. When the modules C3k2_DFF and C2PSA_CAA are integrated—either together or in combination with HDFPN—the model's feature extraction and contextual awareness are substantially enhanced. This leads to a marked increase in R (up to 87.1%), F1, and mAP50; as the model becomes more sensitive to less typical, more ambiguous disease patterns appear. At the same time, a slight decrease in P (from 93.3% to 91.6%) is observed, which is a common phenomenon in multi-module integration: as the detection threshold for difficult cases is lowered, a small number of additional false positives may occur. Nevertheless, the net gain in R and overall performance outweighs this modest P reduction, as reflected in the best F1 and mAP scores for the joint configuration. In summary, these ablation results validate that while HDFPN alone excels in P by minimizing false positives, the synergistic integration of all three modules allows the model to capture more challenging true positives, leading to a more robust and practical system for real-world HLB detection. This trade-off is particularly valuable in complex agricultural environments, where maximizing the detection of difficult or early-stage cases is often more important than achieving the absolute highest P.

To further visually demonstrate the effectiveness of the improved model, a heatmap visualization comparison was conducted on the feature extraction capabilities, as shown in Figure 12. In this comparison, (a) shows the original image, (b) shows the results from YOLO11n, and (c) shows the results from DCH-YOLO11. The heatmaps clearly show that, compared to the original YOLO11 model, the improved DCH-YOLO11 focuses more precisely and intensively on diseased regions, effectively reducing background interference. This indicates that the proposed modules effectively enhance the model's feature discriminability. The combined results from the feature heatmaps and the quantitative ablation analysis thoroughly validate the effectiveness of the modules proposed in this study, as well as the significant improvement and application potential of the overall model.




Figure 12. Comparison of feature extraction heatmaps for different models. (a) Original input image; (b) feature heatmap generated by the original YOLO11n model; (c) feature heatmap generated by the improved DCH-YOLO11 model.

4.4. Generalization Experiments

To evaluate the robustness and generalization capability of the proposed DCH-YOLO11 model under varying data conditions, extensive validation experiments were conducted on the publicly available citrus HLB field symptom recognition dataset constructed by Chi et al. [45]. Unlike the self-constructed MS-HLBD dataset used for training and internal validation, the public dataset provides distinct environmental contexts and varying data distributions, posing a significant challenge for cross-domain evaluation. Therefore, performance on this dataset serves as an essential indicator of the model's adaptability and practical applicability in real-world agricultural scenarios.

The detailed comparative results of DCH-YOLO11 and state-of-the-art object detection models—including the transformer-based RT-DETR and several YOLO variants—are presented in Table 6. As shown, DCH-YOLO11 achieves the highest P (82.7%), F1 (82.2), mAP50 (89.4%), and mAP50-95 (82.6%) among all compared models. Its R (81.8%) is marginally lower than YOLO11n (84.1%), but the overall superior F1 demonstrates a more optimal balance between P and R, underscoring DCH-YOLO11's robust generalization performance. Specifically, compared to the strongest YOLO baselines, DCH-YOLO11 improves P by 1.8% to 5.0% and outperforms RT-DETR by 1.8%. For mAP50, DCH-YOLO11 surpasses RT-DETR by 8.9 percentage points and shows clear gains over the latest YOLO variants—including YOLOv7-tiny, YOLOv8n, YOLOv9-tiny, YOLOv10n, YOLO11n, and YOLOv12n—by 4.0%, 3.8%, 3.2%, 4.7%, 3.2%, and 3.4%, respectively. Its mAP50-95 score of 82.6% further confirms the model's strong performance under stringent IoU thresholds. Importantly, these advancements are achieved with only 2.99M parameters—almost an order of magnitude fewer than RT-DETR (31.99M) and closely matching other lightweight YOLO models—highlighting DCH-YOLO11's balance of accuracy, generalization, and computational efficiency for practical field deployment.

The outstanding generalization performance of DCH-YOLO11 primarily benefits from its specialized module designs. To investigate this further, comprehensive ablation experiments were performed on the same public dataset to quantitatively evaluate the contributions of individual modules, as shown in Table 7. The experimental results indicate that after the addition of the C3k2_DFF module alone, the R reached its highest value of 85.9%, while the F1 improved to 82.5 and the mAP50 metric reached 87.5%. These values are 1.8% and 1.3% higher than the baseline, indicating the strong generalization capability of the C3k2_DFF module on the public dataset. When the C2PSA_CAA module was added

individually, P was 79.7%, R was 85.0%, and mAP50 reached 86.9%. Although the improvements were slightly lower than those of the C3k2_DFF module, the C2PSA_CAA results still showed good generalization, with the lowest parameter count—highlighting its ability to optimize performance while reducing complexity. The addition of the HDFPN module also led to improvements in performance metrics, further confirming the effectiveness of the module. Finally, after integrating all three modules (C3k2_DFF, C2PSA_CAA, and HDFPN) into the complete DCH-YOLO11 model, the model exhibited the best overall performance on the public dataset. Specifically, P reached 82.7%, mAP50 reached 89.4%, and mAP50-95 reached 82.6%, representing improvements of 3.9%, 3.2%, and 3.1%, respectively, over the baseline model. These results underscore the significant generalization advantage of the DCH-YOLO11 model.

Table 6. Performance comparison on the public citrus HLB dataset (Chi et al. [45]). The best results are highlighted in bold.

Model	P/%	R/%	F1	mAP50/%	mAP50-95/%	Parameters
RT-DETR	80.9	82.3	81.6	80.5	73.4	31,998,125
YOLOv7-tiny	79.4	80.7	80.0	85.4	76.4	6,023,832
YOLOv8n	79.2	83.6	81.3	85.6	79.2	3,007,013
YOLOv9-tiny	80.9	81.3	81.1	86.2	80.2	2,619,290
YOLOv10n	77.7	80.7	79.2	84.7	78.0	2,697,146
YOLO11n	78.8	84.1	81.4	86.2	79.5	2,583,517
YOLOv12n	80.0	80.9	80.4	86.0	80.3	2,558,093
DCH-YOLO11	82.7	81.8	82.2	89.4	82.6	2,997,209

Table 7. Ablation results on the public HLB dataset (citrus symptom dataset by Chi et al. [45]) for generalization. “✓” indicates that the module is used, while “-” denotes that the module is not used. Best performance is highlighted in bold.

C3k2_DFF	C2PSA_CAA	HDFPN	P/%	R/%	F1	mAP50 /%	mAP50 -95/%	Parameters
-	-	-	78.8	84.1	81.4	86.2	79.5	2,583,517
✓	-	-	79.3	85.9	82.5	87.5	80.8	2,779,713
-	✓	-	79.7	85.0	82.3	86.9	80.4	2,569,949
-	-	✓	78.5	85.5	81.9	87.1	81.0	2,814,581
✓	✓	✓	82.7	81.8	82.2	89.4	82.6	2,997,209

In summary, the comprehensive evaluations clearly establish that the proposed DCH-YOLO11 model achieves outstanding generalization performance and robustness on external datasets. Not only does it significantly outperform other state-of-the-art detection models, but it also maintains an optimal balance between accuracy, efficiency, and scalability. The ablation analyses further confirm the essential roles and complementary nature of the introduced modules—C3k2_DFF, C2PSA_CAA, and HDFPN—in consistently improving detection accuracy, adaptability, and generalization capabilities. Collectively, these findings demonstrate the strong applicability of DCH-YOLO11 in diverse and practical agricultural environments, particularly in complex real-world HLB detection tasks.

5. Discussion

Although this study utilized a multi-center image dataset spanning various regions, lighting conditions, shooting angles, seasons, and time periods, its overall diversity remains limited. For instance, the sample collection was limited to specific areas, making it difficult to comprehensively capture all variations of HLB symptoms. This limitation primarily stems from constraints in seasonality, geographic accessibility, and labor costs, which led to a relatively concentrated sampling area. Therefore, future research should aim to expand

the data collection scope and increase the sample size from different regions to enhance the representativeness and generalization capability of the dataset.

The proposed DCH-YOLO11 model demonstrated outstanding performance in detection accuracy and other evaluation metrics and also exhibited promising potential for real-world deployment. Specifically, DCH-YOLO11 is designed with a lightweight architecture, containing only 2,995,667 parameters, and achieves a real-time inference speed of 58.4 FPS on an NVIDIA RTX 4080 SUPER GPU, which fully meets the requirements of field-level citrus disease detection. The model is compatible with mainstream deep learning frameworks such as PyTorch, TensorRT, and ONNX, enabling convenient integration into edge devices and agricultural monitoring systems. In addition, all training and evaluation data in this study were collected from real citrus orchards under diverse field conditions, ensuring the model's adaptability to practical deployment scenarios.

Nevertheless, it should be noted that when DCH-YOLO11 is deployed on more resource-constrained devices (such as low-power edge devices or agricultural robots), there may still be trade-offs between maintaining high detection accuracy and achieving real-time performance. To further optimize deployment efficiency, future research could explore advanced model lightweighting techniques, including network pruning, parameter quantization, and knowledge distillation, to achieve even better performance on embedded systems. Additionally, semi-supervised or transfer learning methods could be applied to rapidly enhance the model's generalization using limited new data collected from diverse environments.

6. Conclusions

This study proposes an improved YOLO11-based method for detecting HLB in citrus under natural environmental conditions, addressing issues such as insufficient feature extraction for early subtle disease lesions, low detection accuracy, and poor generalization performance in complex natural environments. The study successfully constructed the MS-HLBD dataset, comprising 9219 images collected and augmented from real-world citrus scenarios. This dataset provides reliable data support for model training and evaluation. Based on the YOLO11 model, this study introduces the DCH-YOLO11 model, achieving significant performance improvements through the innovation of three key modules. Specifically, the designed C3k2_DFF module enhances the model's ability to capture and fuse early disease lesion features, improving the detection of subtle disease characteristics. The C2PSA_CAA module, which incorporates the CAA mechanism, significantly improves the model's capability to recognize fine features in complex leaf vein regions, thereby enhancing detection robustness in challenging backgrounds. Additionally, the HDFPN module optimizes the multi-scale feature fusion strategy, enabling better interaction of features across different scales and effectively improving the model's detection accuracy for objects of various sizes and reducing background interference.

The experimental results show that the improved DCH-YOLO11 model has achieved P of 91.6%, R of 87.1%, F1 of 89.3, mAP50 of 93.1%, and mAP50-95 of 91.5% on the self-constructed dataset. Specifically, mAP50 outperforms models like Faster-RCNN, SSD, RT-DETR, and various YOLO versions (YOLOv7-tiny, YOLOv8n, YOLOv9-tiny, YOLOv10n, YOLO11n, and YOLOv12n) by 13.6%, 8.8%, 5.3%, 3.2%, 2.0%, 1.6%, 2.6%, 1.8%, and 1.6%, respectively. Furthermore, extensive evaluations on both the self-constructed and public citrus HLB datasets demonstrate that DCH-YOLO11 consistently achieves state-of-the-art performance across diverse domains. Notably, on the public dataset, DCH-YOLO11 achieves P of 82.7%, R of 81.8%, F1 of 82.2, mAP50 of 89.4%, and mAP50-95 of 82.6%, all of which are the highest among the compared models. Specifically, mAP50 outperforms RT-DETR, YOLOv7-tiny, YOLOv8n, YOLOv9-tiny, YOLOv10n, YOLO11n, and YOLOv12n

by 8.9%, 4.0%, 3.8%, 3.2%, 4.7%, 3.2%, and 3.4%, respectively. These findings thoroughly validate the effectiveness and versatility of the proposed improvements and module designs under various data conditions.

In summary, the DCH-YOLO11 model provides an efficient and reliable detection method for the precise recognition and control of citrus HLB. Future work may focus on optimizing the model architecture for real-time deployment on resource-constrained devices and enhancing generalization and practical applicability through expanded data collection and transfer learning strategies.

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

Funding: This paper was supported partly by Special projects in key areas for ordinary colleges and universities in Guangdong Province under Grant 2023ZDZX4014, Guangdong Provincial Graduate Education Innovation Plan Project under Grants 2022XSLT056 and 2024JGXM_093, Science and Technology Planning Project of Yunfu under Grant 2023020203, Open Competition Program of Top Ten Critical Priorities of Agricultural Science and Technology Innovation for the 14th Five-Year Plan of Guangdong Province (2022SDZG06, 2023SDZG06, and 2024KJ29), Innovation Team Project of Universities in Guangdong Province under Grant 2021KCXTD019, and Guangzhou Science and Technology Program 201903010043.

Data Availability Statement: All key experimental methods and data analysis procedures are described in detail within the manuscript. The project's code and model implementation are publicly available at <https://github.com/CdW8/DCH-YOLO11> (accessed on 6 July 2025). Further data or materials are available from the corresponding author upon reasonable academic request.

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

Abbreviations

A	Accuracy
AP	Average Precision
AvgPool	Average Pooling
CAA	Context Anchor Attention
CA	Channel Attention
CNN	Convolutional Neural Network
DFF	Dynamic Feature Fusion
DWConv	Depthwise Convolution
F1	F1 score
FLOPs	Floating Point Operations
FPS	Frames Per Second
FN	False Negative
FP	False Positive
FPN	Feature Pyramid Network
HDFPN	High-efficiency Dynamic Feature Pyramid Network
HI	Healthy
HLB	Huanglongbing
HLB_bm	HLB blotchy mottling
HLB_y	HLB yellowing
HLB_Zd	HLB Zinc deficiency
HSFPN	High-Level Screening-feature Fusion Pyramid Network

MaxPool	Maximum Pooling
MS-HLBD	Multi-Symptom HLB Leaf Dataset
mAP	Mean Average Precision
P	Precision
PCR	Polymerase Chain Reaction
qPCR	Real-Time Fluorescence Quantitative PCR
R	Recall
SFF	Select Feature Fusion
SVM	Support Vector Machines
TN	True Negative
TP	True Positive
YOLO	You Only Look Once

References

- Huang, C.; Chen, L.; Xiu, B.; Sun, L.; Gao, J.; Wang, Z.; Ye, F.; Mubasher, H.; Shen, B.; Qiu, D.; et al. Mass Production Technology of *Tamarixia radiata*, Predominant Parasitoid of *Diaphorina citri*, Based on *Citrus reticulata* Blanco. *Chin. J. Biol. Control* **2022**, *38*, 791–796.
- Limayem, A.; Martin, E.M.; Shankar, S. Study on the citrus greening disease: Current challenges and novel therapies. *Microb. Pathog.* **2024**, *192*, 106688. [[CrossRef](#)] [[PubMed](#)]
- Wang, T.; Yang, C.; Yan, Y.; Wang, Y.; Li, R. Cloning and Expression of a Thioredoxin Gene CsTRXh1 from *Citrus sinensis*. *Fujian J. Agric. Sci.* **2022**, *37*, 880–885.
- Wu, F.N.; Xie, S.J.; Zhao, Y.X.; Xu, P.B.; Li, J.J.; Xiao, C.J.; Lin, X.N.; Yang, S.H.; Cen, Y.J.; Huang, J.J. Effects of host plant species on the infection ability of “*Candidatus Liberibacter asiaticus*” and adaptability of *Diaphorina citri*. *J. Environ. Entomol.* **2023**, *45*, 63–72.
- Thakuria, D.; Chaliha, C.; Dutta, P.; Sinha, S.; Uzir, P.; Singh, S.B.; Hazarika, S.; Sahoo, L.; Kharbikar, L.; Singh, D. Citrus Huanglongbing (HLB): Diagnostic and management options. *Physiol. Mol. Plant Pathol.* **2023**, *125*, 102016. [[CrossRef](#)]
- Zhang, T.; Yang, G.; He, X.; Liu, Q.; Li, D. Research progress on detection and control method of citrus Huanglongbing. *Sci. Technol. Rev.* **2024**, *42*, 75–83. (In Chinese)
- Xu, Q.; Su, Y.; Sun, L.; Cai, J. Detection of citrus Huanglongbing at different stages of infection using a homemade electronic nose system. *Comput. Electron. Agric.* **2025**, *229*, 109845. [[CrossRef](#)]
- Keremane, M.L.; McCollum, T.G.; Roose, M.L.; Lee, R.F.; Ramadugu, C. An improved reference gene for detection of “*Candidatus Liberibacter asiaticus*” associated with citrus huanglongbing by qPCR and digital droplet PCR assays. *Plants* **2021**, *10*, 2111. [[CrossRef](#)]
- Frederick, Q.; Burks, T.; Yadav, P.K.; Qin, J.; Kim, M.; Dewdney, M. Classifying adaxial and abaxial sides of diseased citrus leaves with selected hyperspectral bands and YOLOv8. *Smart Agric. Technol.* **2024**, *9*, 100600. [[CrossRef](#)]
- Ye, N.; Mai, W.; Qin, F.; Yuan, S.; Liu, B.; Li, Z.; Liu, C.; Wan, F.; Qian, W.; Wu, Z.; et al. Early detection of Citrus Huanglongbing by UAV remote sensing based on MGA-UNet. *Front. Plant Sci.* **2025**, *16*, 1503645. [[CrossRef](#)]
- Xu, Q.; Cai, J.R.; Zhang, W.; Bai, J.W.; Li, Z.Q.; Tan, B.; Sun, L. Detection of citrus Huanglongbing (HLB) based on the HLB-induced leaf starch accumulation using a home-made computer vision system. *Biosyst. Eng.* **2022**, *218*, 163–174. [[CrossRef](#)]
- Liu, Y.; Xiao, H.; Xu, H.; Rao, Y.; Jiang, X.; Sun, X. Visual discrimination of citrus HLB based on image features. *Vib. Spectrosc.* **2019**, *102*, 103–111. [[CrossRef](#)]
- Yan, K.; Song, X.; Yang, J.; Xiao, J.; Xu, X.; Guo, J.; Zhu, H.; Lan, Y.; Zhang, Y. Citrus huanglongbing detection: A hyperspectral data-driven model integrating feature band selection with machine learning algorithms. *Crop Prot.* **2025**, *188*, 107008. [[CrossRef](#)]
- Aswini, E.; Vijayakumaran, C. Auto Detector for Huanglongbing Citrus Greening Disease using YOLOV7. In Proceedings of the 2023 World Conference on Communication & Computing (WCONF), Raipur, India, 14–16 July 2023; IEEE: Piscataway, NJ, USA, 2023; pp. 1–6.
- Li, S.; Liang, Q.; Yu, Y.; Chen, Y.; Fu, H.; Zhang, H. Research on Asian citrus psyllid YOLO v8-MC recognition algorithm and insect remote monitoring system. *Trans. Chin. Soc. Agric. Mach.* **2024**, *55*, 210–218. (In Chinese)
- Lin, Y.; Liang, J.; Liu, S.; Jia, S.; Yu, J.; Hou, Y. Individual identification of citrus Huanglongbing based on UAV visible-light imagery and support vector machine. *Southwest China J. Agric. Sci.* **2022**, *35*, 2554–2563. (In Chinese)
- Lu, J.; Lin, J.; Huang, Z.; Wang, W.; Qiu, H.; Yang, R.; Chen, P. Identification of citrus fruit infected with Huanglongbing based on Mixup algorithm and convolutional neural network. *J. South China Agric. Univ.* **2021**, *42*, 94–101. (In Chinese)
- Hakim, A.; Srivastava, A.K.; Hamza, A.; Owais, M.; Habib-ur Rahman, M.; Qadri, S.; Qayyum, M.A.; Ahmad Khan, F.Z.; Mahmood, M.T.; Gaiser, T. Yolo-pest: An optimized YoloV8x for detection of small insect pests using smart traps. *Sci. Rep.* **2025**, *15*, 14029. [[CrossRef](#)]

19. Silva, C.E.S.E.; Fragoso, J.B.; Paixão, T.; Alvarez, A.B.; Palomino-Quispe, F. A Low Computational Cost Deep Learning Approach for Localization and Classification of Diseases and Pests in Coffee Leaves. *IEEE Access* **2025**, *13*, 71943–71964. [\[CrossRef\]](#)
20. Jin, S.; Zhou, L.; Zhou, H. CO-YOLO: A lightweight and efficient model for *Camellia oleifera* fruit object detection and posture determination. *Comput. Electron. Agric.* **2025**, *235*, 110394. [\[CrossRef\]](#)
21. Zhu, F.; Zhang, W.; Li, Z.; Gao, T.; Zhao, Q. A YOLOv7-based robotic harvesting system for *Agaricus bisporus* using a depth camera. *Turk. J. Agric. For.* **2025**, *49*, 380–396. [\[CrossRef\]](#)
22. Goyal, R.; Nath, A.; Niranjan, U.; Sharda, R. Analyzing the performance of deep convolutional neural network models for weed identification in potato fields. *Crop Prot.* **2025**, *188*, 107035. [\[CrossRef\]](#)
23. Shi, H.; Liu, C.; Wu, M.; Zhang, H.; Song, H.; Sun, H.; Li, Y.; Hu, J. Real-time detection of Chinese cabbage seedlings in the field based on YOLO11-CGB. *Front. Plant Sci.* **2025**, *16*, 1558378. [\[CrossRef\]](#)
24. Niu, Z.; Zhao, Y. A Tea Buds Detection Algorithm Based on Improved YOLO11. In Proceedings of the 2024 International Conference on Image Processing, Computer Vision and Machine Learning (ICICML), Shenzhen, China, 22–24 November 2024; IEEE: Piscataway, NJ, USA, 2024; pp. 1550–1553.
25. Li, P.; Chen, J.; Chen, Q.; Huang, L.; Jiang, Z.; Hua, W.; Li, Y. Detection and picking point localization of grape bunches and stems based on oriented bounding box. *Comput. Electron. Agric.* **2025**, *233*, 110168. [\[CrossRef\]](#)
26. Bassanezi, R.B.; Belasque, J., Jr.; Montesino, L. Frequency of symptomatic trees removal in small citrus blocks on citrus huanglongbing epidemics. *Crop Prot.* **2013**, *52*, 72–77. [\[CrossRef\]](#)
27. Kwakye, S.; Kadyampakeni, D.M. Micronutrients improve growth and development of HLB-affected citrus trees in Florida. *Plants* **2022**, *12*, 73. [\[CrossRef\]](#)
28. Anzo Hernández, A.; Giménez Mujica, U.J.; Hernández Gracidas, C.A.; Oliveros Oliveros, J.J. Optimizing control parameters for Huanglongbing disease in citrus orchards using SAIR-SI compartmental model, epidemic final size, and genetic algorithms. *J. Math. Biol.* **2025**, *90*, 4. [\[CrossRef\]](#)
29. Silva, L.M.; Martins, E.C.; Ferreira, A.A.P.; Wulff, N.A.; Yamanaka, H. Impedimetric immunosensor versus qPCR for Huanglongbing detection. *Talanta* **2025**, *283*, 127132. [\[CrossRef\]](#)
30. Wang, X.; Wu, Y.; Cui, L.; Qian, H.; Li, B.; Wang, X. Linear pattern detection of building groups by integrating dynamic snake convolution with YOLO11. *Geocarto Int.* **2025**, *40*, 2471914. [\[CrossRef\]](#)
31. Wei, J.; Ni, L.; Luo, L.; Chen, M.; You, M.; Sun, Y.; Hu, T. GFS-YOLO11: A Maturity Detection Model for Multi-Variety Tomato. *Agronomy* **2024**, *14*, 2644. [\[CrossRef\]](#)
32. Sapkota, R.; Meng, Z.; Karkee, M. Synthetic meets authentic: Leveraging llm generated datasets for yolo11 and yolov10-based apple detection through machine vision sensors. *Smart Agric. Technol.* **2024**, *9*, 100614. [\[CrossRef\]](#)
33. Yang, J.; Qiu, P.; Zhang, Y.; Marcus, D.S.; Sotiras, A. D-net: Dynamic large kernel with dynamic feature fusion for volumetric medical image segmentation. *arXiv* **2024**, arXiv:2403.10674.
34. Cai, X.; Lai, Q.; Wang, Y.; Wang, W.; Sun, Z.; Yao, Y. Poly Kernel Inception Network for Remote Sensing Detection. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Seattle, WA, USA, 16–22 June 2024; pp. 27706–27716.
35. Chen, Y.; Zhang, C.; Chen, B.; Huang, Y.; Sun, Y.; Wang, C.; Fu, X.; Dai, Y.; Qin, F.; Peng, Y.; et al. Accurate leukocyte detection based on deformable-DETR and multi-level feature fusion for aiding diagnosis of blood diseases. *Comput. Biol. Med.* **2024**, *170*, 107917. [\[CrossRef\]](#) [\[PubMed\]](#)
36. Ren, S.; He, K.; Girshick, R.; Sun, J. Faster r-cnn: Towards real-time object detection with region proposal networks. *Adv. Neural Inf. Process. Syst.* **2015**, *28*, 91–99. [\[CrossRef\]](#)
37. Liu, W.; Anguelov, D.; Erhan, D.; Szegedy, C.; Reed, S.; Fu, C.Y.; Berg, A.C. SSD: Single shot multibox detector. In Proceedings of the Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, 11–14 October 2016; Proceedings, Part I 14; Springer: Cham, Switzerland, 2016; pp. 21–37.
38. Zhao, Y.; Lv, W.; Xu, S.; Wei, J.; Wang, G.; Dang, Q.; Liu, Y.; Chen, J. Detrs beat yolos on real-time object detection. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, 16–22 June 2024; pp. 16965–16974.
39. Wang, C.Y.; Bochkovskiy, A.; Liao, H.Y.M. YOLOv7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Vancouver, BC, Canada, 17–24 June 2023; pp. 7464–7475.
40. Varghese, R.; Sambath, M. Yolov8: A novel object detection algorithm with enhanced performance and robustness. In Proceedings of the 2024 International Conference on Advances in Data Engineering and Intelligent Computing Systems (ADICS), Chennai, India, 18–19 April 2024; IEEE: Piscataway, NJ, USA, 2024; pp. 1–6.
41. Wang, C.Y.; Yeh, I.H.; Mark Liao, H.Y. Yolov9: Learning what you want to learn using programmable gradient information. In Proceedings of the European Conference on Computer Vision, Milan, Italy, 29 September–4 October 2024; Springer: Cham, Switzerland, 2024; pp. 1–21.

42. Wang, A.; Chen, H.; Liu, L.; Chen, K.; Lin, Z.; Han, J.; Ding, G. Yolov10: Real-time end-to-end object detection. *arXiv* **2024**, arXiv:2405.14458.
43. Khanam, R.; Hussain, M. Yolov11: An overview of the key architectural enhancements. *arXiv* **2024**, arXiv:2410.17725.
44. Tian, Y.; Ye, Q.; Doermann, D. Yolov12: Attention-centric real-time object detectors. *arXiv* **2025**, arXiv:2502.12524.
45. Chi, M.; Chen, S.; Huang, T.; Chen, S.; Liang, Y.; Qiu, R. Image datasets of citrus Huanglongbing field symptom recognition. *China Sci. Data* **2025**, *10*, 52–61. (In Chinese)

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.