

Towards Efficient and Intelligent Laser Weeding: Method and Dataset for Weed Stem Detection

Dingning Liu^{1*}, Jinzhe Li^{1*}, Haoyang Su^{1*}, Bei Cui^{2†}, Zhihui Wang³, Qingbo Yuan²,
Wanli Ouyang¹, Nanqing Dong^{1†}

¹Shanghai Artificial Intelligence Laboratory

²Yazhouwan National Laboratory

³Dalian University of Technology

{liudingning, dongnanqing}@pjlab.org.cn

Abstract

Weed control is a critical challenge in modern agriculture, as weeds compete with crops for essential nutrient resources, significantly reducing crop yield and quality. Traditional weed control methods, including chemical and mechanical approaches, have real-life limitations such as associated environmental impact and efficiency. An emerging yet effective approach is laser weeding, which uses a laser beam as the stem cutter. Although there have been studies that use deep learning in weed recognition, its application in intelligent laser weeding still requires a comprehensive understanding. Thus, this study represents the first empirical investigation of weed recognition for laser weeding. To increase the efficiency of laser beam cut and avoid damaging the crops of interest, the laser beam shall be directly aimed at the weed root. Yet, weed stem detection remains an under-explored problem. We integrate the detection of crop and weed with the localization of weed stem into one end-to-end system. To train and validate the proposed system in a real-life scenario, we curate and construct a high-quality weed stem detection dataset with human annotations. The dataset consists of 7,161 high-resolution pictures collected in the field with annotations of 11,151 instances of weed. Experimental results show that the proposed system improves weeding accuracy by 6.7% and reduces energy cost by 32.3% compared to existing weed recognition systems.

Code & Dataset —

<https://github.com/open-sciencelab/WeedStemDetection>

Introduction

Sustainable agricultural management is essential for addressing global hunger and achieving the UN’s “Zero Hunger” goal and the principle of “Leaving No One Behind” (LNOB) (United Nations 2023). Effective weed control plays a non-trivial role in maintaining food security, as weeds compete with crops for critical resources such as water, nutrients, and sunlight, which directly affects crop yield and quality.

*These authors contributed equally.

†Corresponding author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

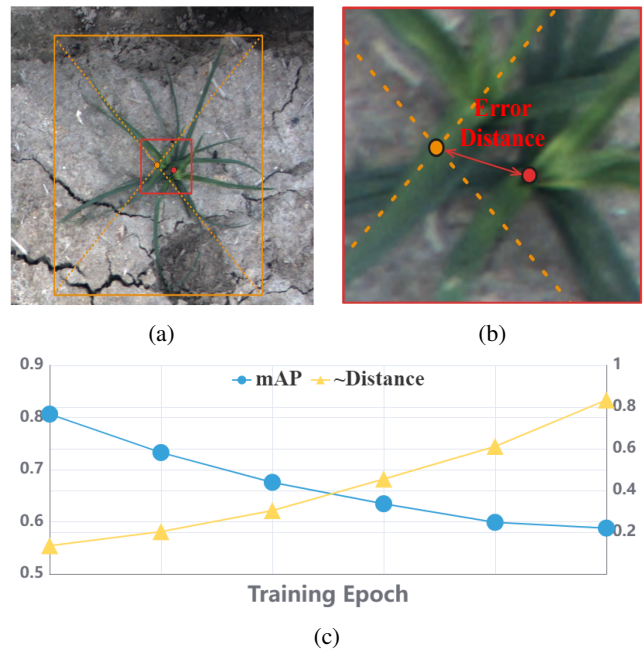


Figure 1: (a) Weed detection with YOLOv7: Yellow dashed lines mark bounding box diagonals, indicating the geometric center, while the blue dot shows the ground-truth weed stem location. (b) Close-up of the red-boxed region, highlighting misalignment between the bounding box center and the ground-truth point. (c) In laser weeding, better weed detection performance (lower mAP) does not always mean better weed stem detection (lower Euclidean Distance).

Current weed control methods are generally categorized into chemical weeding and mechanical weeding. Chemical methods usually use toxic substances to inhibit or destroy weeds at various growth stages, including those applied before or after weeds emerge. Although effective against different weed types, these methods can negatively influence the crop quality and inevitably cause chemical soil degradation, resulting in environmental pollution (Zhang 1996a,b; Tanveer et al. 2003). Instead, mechanical weed control meth-

ods use machines such as mowers (Pirchio et al. 2018; Sportelli et al. 2020; Aamlid et al. 2021). Mowers often miss small weeds and can only cut the part of weed above the surface. Thus, mowers are not effective when dealing with deeply-rooted and perennial ones, leading to frequent maintenance. A desired weed control approach should take both efficiency and environmental friendliness into consideration. Laser weeding (Carbon Robotics 2022; WeLASER 2023) offers a promising alternative with aforementioned properties. It leverages high-energy and high-temperature laser beam to target and cut weeds at the stem, effectively killing the weeds. Additionally, laser weeding can be environmentally friendly if it is powered by clean energy.

Fueled by recent advances in deep learning, weed recognition has been well studied theoretically (Wu et al. 2021; Hu et al. 2024). Considering that laser weeding has the advantages in weeding efficiency and energy consumption, intelligent laser weeding seems to be a promising path for society. On the contrary, efficiency and energy become two critical issues for intelligent laser weeders. To conserve high-energy beams, the diameter of laser transmitters is much smaller in size compared with the leaves. Simply detecting or segmenting the weed is not an effective signal to transmit the laser beam as cutting leaves can not eradicate the weed. To maintain high efficiency in terms of energy usage, the laser beam shall be aimed directly at the bottom of weed stem. Meanwhile, as the laser can cause irreversible damage on crops, the detection algorithms require low false positive rate. So far, though there have been a few start-ups trying to address this challenge (Carbon Robotics 2022; WeLASER 2023), accurately locating the weed stem remains a challenge (Zhang, Zhong, and Zhou 2023). A visual illustration is presented in Fig. 1. Based on geometric principle and agricultural knowledge, weed stems are intuitively expected to align with the geometric center of the detected bounding box in a vertical view. Though the predicted bounding boxes can achieve high mAP, the predicted location of weed stem is far from the ground truth location (Fig. 1(b)). Moreover, traditional weed detection methods typically use mAP as the evaluation metric, which may not be suitable for the task of laser weeding. As shown in Fig. 1(c), a method with a high mAP score may still exhibit poor root localization performance, but improving distance accuracy is crucial as it directly benefits the effectiveness of laser weeding.

To tackle the aforementioned challenges, we propose a pipeline that integrates crop and weed detection with weed stem localization into a unified end-to-end system. Specifically, we introduce an additional root coordinate regression branch within the object detection framework. The proposed system can process a sequence of images or a real-time video stream, detecting plant bounding boxes and simultaneously pinpointing weed stems for laser transmission, thereby ensuring effective weed control without damaging crops. This pipeline is simple yet robust, and can be easily implemented in object detectors. To train and validate the proposed system in real-world conditions, and to empirically understand the task of weed recognition under the setup of laser weeding, we collect and curate the Weed Stem Detection (WSD) dataset, consisting of 7,161 high-resolution im-

ages with 11,151 annotated instances. This dataset includes bounding boxes for three crops and weeds, as well as the coordinates of weed stem. The main contributions of this work are summarized below.

- We provide a high-quality weed stem dataset with human annotations and the first empirical study on weed recognition for practical laser weeding, addressing a significant academic gap in laser weeding.
- We propose an end-to-end deep learning pipeline that integrates crop and weed detection with weed stem localization and can be extended to semi-supervised learning, which can further leverage unlabeled data.
- We experimentally demonstrate that our method is more efficient than previous detection-based methods in laser weeding by improving the weeding accuracy by 6.7% and reducing the energy cost by 32.3%.

Related Work

Weed Datasets

Existing weed datasets primarily address weed recognition or detection (Hasan et al. 2021; Hu et al. 2024). We summarize the key public datasets with human annotations in Tab. 1. DeepWeeds (Olsen et al. 2019) includes 17,509 images of eight Australian weeds but lacks crop data, limiting its practical weeding applications. The Weed-Corn/Lettuce/Radish dataset (Jiang et al. 2020) contains 7,200 images with four species (three crops and one weed), while the Food Crop and Weed Dataset (Sudars et al. 2020) includes 1,118 images of seven species (six crops and one weed). CottonWeedID15 (Chen et al. 2022) consists of 5,187 weed images from cotton fields with image-level annotations only. CottonWeedDet12 (Lu 2023) and CottonWeedDet3 (Rahman, Lu, and Wang 2023) add bounding box annotations. The largest dataset, CropAndWeed (Steininger et al. 2023), has coarse machine-generated labels. However, precise weed stem localization is essential for laser weeding, requiring high-quality annotations in terms of localization. Although some laser-weeding datasets exist (Zhang et al. 2024), they focus on classification, detection, and segmentation rather than precise stem localization. To the best of our knowledge, WSD is the first dataset with human annotations for both crop and weed detection, as well as weed stem localization.

Weed Recognition

Existing studies on weed recognition can be broadly categorized into four tasks: weed classification, weed object detection, weed object segmentation, and weed instance segmentation (Hu et al. 2024). Weed classification focuses on identifying weeds at the image level, determining whether an image contains non-crop plants. For instance, SVM classifiers have achieved about 95% accuracy in relatively simple environments (Zhang et al. 2022), and by combining VGG with SVM (Tao and Wei 2022), a 99% accuracy rate has been reached in distinguishing between weeds and grapevines. Weed object detection extends beyond classification by providing bounding boxes to locate weeds within images. (Parra

et al. 2020; Nasiri et al. 2022) Various models have been successfully applied to this task, including DetectNet (Yu et al. 2019), Faster R-CNN (Veeranampalayam Sivakumar et al. 2020), and YOLOv3 (Sharpe et al. 2020), all showing promising results. Weed object segmentation and instance segmentation focus on pixel-level recognition (Jeon, Tian, and Zhu 2011; Long, Shelhamer, and Darrell 2015; You, Liu, and Lee 2020), offering more detailed analysis. For example, VGG-UNet has been used to segment sugar beets and weeds (Fawakherji et al. 2019). However, none of these methods can localize the weed stem, a crucial aspect of effective weed management. To address this gap, our work introduces an end-to-end framework that simultaneously detects crops and weeds while localizing the weed stem.




Figure 2: Image samples show raw images (left) and 16x zoomed sections (right), highlighting four different species.




Figure 3: Distribution of instance annotations per image: The X-axis shows the number of instances (including both bounding box and point annotations) per image, while the Y-axis indicates the corresponding image count.

Weed Stem Detection Dataset

Data Collection

The standard RGB images are collected by a custom-built autonomous vehicle equipped with Teledyne FLIR BFS-U3-123S6C-C, a high-resolution imagery sensor. Each image has a resolution of 2048×2048 . The sensor is embedded in the autonomous vehicle, making the sensor at a relatively fixed height above the surface, which is one meter for the prototype vehicle. The images are captured in three different experimental fields planted with three different crops: maize, soybean, and mungbean, respectively. We intentionally planted weed seeds in the field at staggered intervals, resulting in weeds at various growth stages. All crops, however, are at the seedling stage, 30 days after sowing—an important period for weeding.

Data Annotation

We deployed LabelImg¹, a graphical image annotation tool. The human annotators can label object bounding boxes in images with LabelImg, which saves the annotation details in XML files. Three professional agronomists with advanced graduate degrees and field experience were hired to complete the annotation. The annotation process was completed in two steps. First, the bounding boxes of crop and weed were annotated. Then, weed stem locations were annotated in a point coordinate fashion. All final annotations were verified by all three agronomists to achieve consensus. Any discrepancies among the human annotators were resolved through re-annotation to ensure reliability. Fig. 2 shows four annotated images with zoomed-in visualization.

Dataset Statistics

There are 7,161 images in total, with 1,556 annotated and 5,605 unannotated images. The inclusion of unannotated images allows the dataset to be extended for semi-supervised learning scenarios. The distribution of instance annotations per image is illustrated in Fig. 3. It is worth mentioning that the annotation is time-consuming. On average, it takes

¹<https://pypi.org/project/labelImg/>

Dataset	Stem	BBox	# Img	# Species	# Inst	Res
DeepWeeds (2019)	✗	✗	17509	1	-	256 × 256
Weed-Corn/Lettuce/Radish (2020)	✗	✗	7200	4	-	800 × 600
Crops and Weed Dataset (2020)	✗	✓	1118	7	-	[480, 1000] × [384, 1280]
CottonWeedID15 (2022)	✗	✗	5187	1	-	512 × 512
CottonWeedDet3 (2023)	✗	✓	848	1	1.8 ± 1.5	4442 × 4335
CottonWeedDet12 (2023)	✗	✓	5648	1	1.7 ± 1.4	3024 × 4032
WSD (Ours)	✓	✓	7161	4	12.5 ± 7.5	2048 × 2048

Table 1: Comparison between WSD dataset and existing weed datasets with human annotations. “Stem” indicates whether stem annotations are provided. “#Img” denotes the number of images. “# Species” denotes the number of species. “# Inst” denotes the average number of annotations per image, along with the standard deviation. “Res” denotes the image resolution.

approximately 135 seconds to label a weed instance. The statistics of WSD are summarized in Tab. 2.

Class	# Instances	Share	# Images	Time (s)
Weed	11151	57.3%	1379	135 ± 45
Maize	5808	29.9%	1175	105 ± 15
Soybean	1624	8.4%	268	60 ± 10
Mungbean	848	4.4%	100	75 ± 15

Table 2: Statistics of Weed Stem Detection Dataset. “Instances” indicates the number of bounding boxes with additional point annotations. “Share” represents the percentage of instances in this category. “Images” is the number of images containing this category. “Time” refers to the average ± standard deviation of annotation time in seconds by professional agronomists.

Method

The proposed laser weeding pipeline is depicted in Fig. 4. The autonomous vehicle first captures images that include both crops and weeds. These images are then processed by a neural network to detect and localize weed stems. Upon detection, a laser beam is emitted to cut the weed stems. This section details the integration of stem regression within a pre-trained object detection neural network and the subsequent enhancement of its performance using semi-supervised learning.

Weed Stem Regression

To accurately localize weed stems, we augment the pre-trained object detection neural network $NN(\cdot)$ with an additional stem coordinate regression head, formulated as:

$$E_i = NN(I_i) \quad (1)$$

$$\hat{y}_i = Conv(E_i) \quad (2)$$

where I_i represents the i -th input image, E_i denotes the extracted image embedding, \hat{y}_i is the predicted stem coordi-

nate, and $Conv(\cdot)$ is the regression head. The regression loss L_{reg} is computed as:

$$L_{reg} = \frac{1}{n} \sum_{i=1}^n MSE(y_i, \hat{y}_i) \quad (3)$$

$MSE(\cdot)$ represents the Euclidean Distance calculation, y_i is the ground truth weed stem coordinate, and only weed coordinates are used in calculating the regression loss. To jointly optimize bounding box detection and weed stem regression, we combine the regression loss L_{reg} with the classification loss L_{cls} and the bounding box detection loss L_{bbox} as follows:

$$L = \alpha \cdot L_{cls} + \beta \cdot L_{bbox} + \gamma \cdot L_{reg} \quad (4)$$

where α , β , and γ are hyper-parameters that balance the contributions of the different losses.

Extension to Leverage Unlabelled Images

Labeling images in real-world scenarios is labor-intensive. To reduce annotation costs and simultaneously leverage unlabeled images to enhance model performance, we employ a teacher-student framework for semi-supervised learning (Kingma et al. 2014; Zhai et al. 2019; Berthelot et al. 2019; Xu et al. 2021). As illustrated in Fig. 5, pseudo labels for the unlabeled data \mathcal{D}_u are generated using a teacher model. The student model is then trained on both the labeled data \mathcal{D}_l and the pseudo-labeled data \mathcal{D}_p .

Pseudo Label Generation To effectively utilize the abundant unlabeled images, we first fine-tune a teacher model $Teacher(\cdot)$ based on the pre-trained neural network with the combined loss L . The fine-tuned teacher model classifies unlabeled images, assigning pseudo labels to predictions with confidence higher than the threshold τ . Confidence $ConfScore$ is calculated as:

$$ConfScore = Max(Softmax(Teacher(E_u))) \quad (5)$$

where E_u denotes unlabeled image embeddings (subscripts are omitted for simplicity). Since precise weed coordinate prediction is crucial, we use ground-truth weed image embeddings as anchors to filter out low-quality predictions.

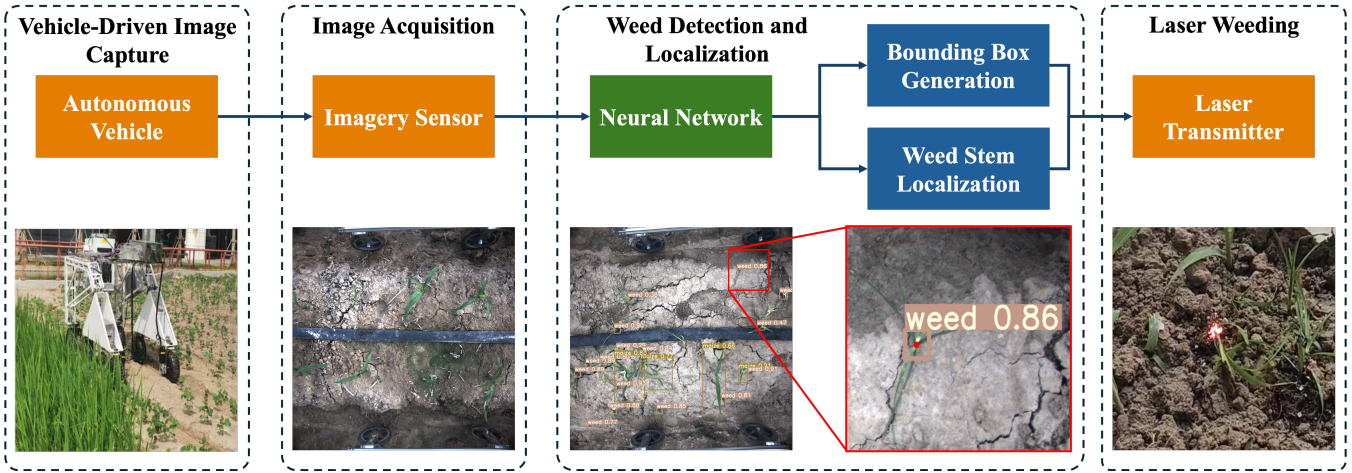


Figure 4: The pipeline of intelligent laser weeding. The autonomous vehicle captures the image (or video frame). The proposed neural network infers the class and location of each crop and weed. Upon identifying a weed, the model also outputs the stem location of the weed, followed by a laser beam cut.

Specifically, we extract weed embeddings E_l^w from labeled data and store them in a weed bank. The cosine similarity between predicted weed embeddings E_u^w and all pre-extracted weed embeddings E_l^w is then calculated:

$$SimScore = CosineSimilarity(E_l^w, E_u^w) \quad (6)$$

Predictions with $SimScore$ higher than the threshold ξ are assigned as pseudo labels. Finally, the student model is trained on both labeled data \mathcal{D}_l and pseudo-labeled data \mathcal{D}_p . Notably, weak and strong augmentations are applied to each unlabeled image: the weakly augmented images are fed into the student network, while the strongly augmented images are processed by the teacher network. Weak augmentations include adjustments to brightness and contrast, while strong augmentations additionally involve cropping and flipping.

Model	Dist↓	FP	Accuracy	Cost
Detection	2.9770	0	75.37%	1.55
Ours	2.4838	0	80.42%	1.05

Table 3: Effect of stem regression. “FP” denotes false positive rate, *i.e.* the crops are identified as the weeds. “Accuracy” denotes the weeding accuracy. “Cost” denotes the energy cost where the unit cost denotes the hypothetical minimum energy cost with ground truth-level prediction, *i.e.* one weed only requires one shot.

Experiments

Experimental Setup

Implementation We conduct empirical evaluations of our method on the WSD dataset, using 80% of the data for training, 10% for validation, and 10% for testing. Given the hardware limitations of the autonomous vehicle, we prioritize lightweight deployment and real-time inference by selecting

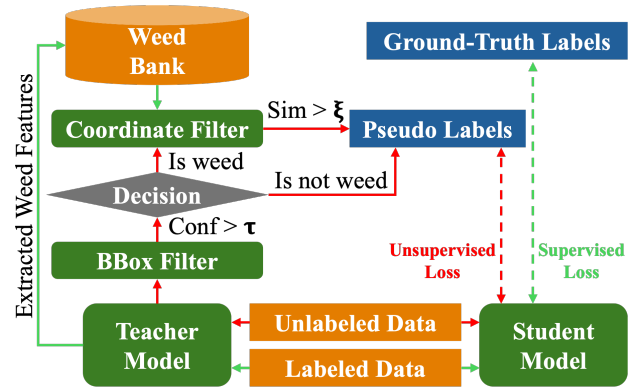


Figure 5: Overview of semi-supervised learning process. Pseudo labels are first generated for unlabeled data using a teacher model, followed by training a student model with both labeled and pseudo-labeled data. “Conf” represents the confidence score for classification, and “Sim” denotes the cosine similarity between extracted ground-truth weed embeddings and predicted weed embeddings, used to filter out low-quality weed localization. τ and ξ are hyper-parameters.

YOLOv7 (Wang, Bochkovskiy, and Liao 2023) as the baseline detection model due to its superior performance among existing object detection methods (Dang et al. 2023; Rahman, Lu, and Wang 2023). All models are trained using an SGD optimizer with a learning rate of $1e-3$ for 300 epochs. We set the loss weights α , β , and γ to 0.2, 0.3, and 0.5, respectively. For semi-supervised learning, we adopt a dynamic updating mechanism with EMA, where the smoothing factor is set to 0.9, the confidence threshold is set to 0.5 and the cosine similarity threshold to 0.4. To be mentioned, all hyper parameters are determined on the basis of empirical results. The experiments are conducted on a single NVIDIA Tesla A100 80G GPU.




Figure 6: Qualitative comparison of weed detection results between the vanilla detection method and our method. The results of our method (second row) shows higher true positive rate than the vanilla detection method (first row), especially when a weed spread and looks like multiple weeds, thus conserving the energy.

Evaluation Metric While mAP is a well-established metric that aggregates precision and recall to provide an overall measure of detection performance, it primarily evaluates the presence and classification of objects rather than their precise positioning. As shown in Fig. 1, mAP is not a robust measure for the task of interest. Instead, we evaluate using Euclidean Distance (Dist), where a lower value indicates higher accuracy in weed stem localization.

Results

Effect of Stem Regression In this empirical study, we first validate the effectiveness of stem regression, which significantly reduces the Dist value (Tab. 3). We then design a real-world simulated experiment to assess the efficiency of the proposed system. Under identical experimental conditions, the vanilla detection method and our method process the same number of weeds independently. For the vanilla method, the geometric center of the predicted bounding box is used as the stem coordinate. We evaluate two aspects: weeding accuracy (percentage of weeds eradicated) and energy consumption (laser shots). As shown in Tab. 3, our method reduces energy cost by up to 32.3% while improving accuracy by 6.7%, with no misidentification of crops as weeds. This suggests a clear visual difference between crops and weeds. Qualitative comparisons of weed detection and stem localization are shown in Fig. 6 and Fig. 7.

Necessity of Object Detection With WSD, another learning choice is to regress on weed stem coordinates directly, without the object detection framework. As shown in Tab. 4, the integration of stem regression and weed detection is im-


portant for the task of interest. Specifically, adding weed detection reduces the distance error from 4.3306 to 2.4838, demonstrating a significant performance gain.

Stem Reg.	Det.	Unlabeled	Dist↓
✓			4.3306
✓	✓		2.4838
✓	✓	✓	2.1485

Table 4: Analysis on learning components. “Stem Reg.” denotes stem regression. “Det.” denotes standard object detection. “Unlabeled” denotes unlabeled data, which means semi-supervised learning.

Study on Semi-Supervised Learning Using the same setup as above, we evaluated the effectiveness of our semi-supervised learning method, incorporating unlabeled images during training. With semi-supervised learning, the Dist value was further reduced to 2.1485. Additionally, we replace the detection backbone from YOLOv7 to YOLOv8 (Reis et al. 2023). Fig. 8 suggests that the extension to semi-supervised learning is model agnostic.

Ablation Study on Detection Backbone Although YOLOv8 generally outperforms YOLOv7 in object detection, it shows lower performance in Fig. 8. We hypothesize that the detection backbone plays a role in this discrepancy. Recent studies (Joiya 2022; Nanos 2023; Key-labs 2023; Roboflow 2023) indicate that single-stage mod-



(a) Detection

(b) Ours

Figure 7: Qualitative comparison of stem localization results between the vanilla detection method (a) and our method (b). Green and red points denote the prediction and the ground truth, respectively. Our method consistently outperforms the vanilla detection method.

els, particularly those in the YOLO family, excel in tasks requiring both high-speed inference and competitive accuracy, compared to Faster-RCNN (Ren et al. 2016) and SSD (Liu et al. 2016). We evaluate YOLOv7, YOLOv8, and YOLOv10 (Wang et al. 2024), which are anchor-based, anchor-free, and anchor-free, respectively. A comparative analysis is shown in Tab. 5.

With stem prediction incorporated, all models achieved a 0% False Positive rate, ensuring no seedlings were misidentified as weeds. This confirms the effectiveness of our approach in reducing crop damage risk during automated weeding, enhancing both accuracy and safety for young plants. Despite expectations that YOLOv8 and YOLOv10 would perform better, their anchor-free approach resulted in significant drift in stem predictions. In contrast, YOLOv7’s anchor-based method prevents drift by keeping loss calculation confined to the respective grid. The results in Tab. 5 further validate YOLOv7 as the baseline detection model,

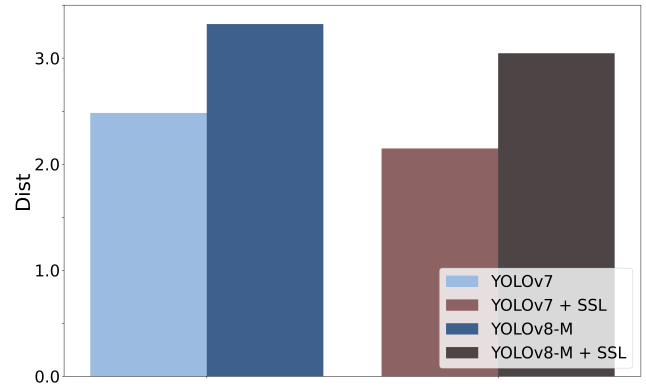


Figure 8: Effect of semi-supervised learning. Semi-supervised learning consistently improves the model performance. “SSL” denotes semi-supervised learning.

balancing computational efficiency and accuracy.

Model	#Param ↓	Time(ms) ↓	FP ↓	Dist ↓
YOLOv10-M	26M	130.4	0	3.2691
YOLOv8-M	26M	407.3	0	3.3219
YOLOv8-L	43M	597.2	0	3.1473
YOLOv7-X	142M	119.5	0	2.4426
YOLOv7	37M	60.2	0	2.4838

Table 5: Performance comparison among different YOLO backbones. “#Param” denotes the number of parameters. “Time(ms)” denotes the inference time. “FP” denotes the false positive rate.

Effect of the Classifier Threshold We set a confidence threshold of 0.15 during inference, consistent with most zero-shot object detection models. As shown in Tab. 6, the minimum zero-FP threshold is 0.056.

Threshold	FP
0.150	0.000
0.056	0.000
0.055	0.005
0.050	0.020

Table 6: Effect of the confidence threshold: raising the threshold can enhance generalization in zero-shot scenarios.

Conclusion

In this work, we propose an end-to-end pipeline that unifies crop and weed detection and weed stem localization, which shows improved accuracy and reduced energy cost. This study not only serves an empirical study of practical weed recognition, but also poses a promising research direction on intelligent laser weeding.

Acknowledgments

This work is supported by Shanghai Artificial Intelligence Laboratory.

References

- Aamlid, T.; Hesselsøe, K. J.; Pettersen, T.; and Borchert, A. F. 2021. ROBO-GOLF: Robotic mowers for better turf quality on golf course fairways and semi-roughs, Results from 2020. *NIBIO Rapport*.
- Berthelot, D.; Carlini, N.; Goodfellow, I.; Papernot, N.; Oliver, A.; and Raffel, C. A. 2019. Mixmatch: A holistic approach to semi-supervised learning. *Advances in neural information processing systems*, 32.
- Carbon Robotics. 2022. Laserweeder Implement.
- Chen, D.; Lu, Y.; Li, Z.; and Young, S. 2022. Performance evaluation of deep transfer learning on multi-class identification of common weed species in cotton production systems. *Computers and Electronics in Agriculture*, 198: 107091.
- Dang, F.; Chen, D.; Lu, Y.; and Li, Z. 2023. YOLOWeeds: A novel benchmark of YOLO object detectors for multi-class weed detection in cotton production systems. *Computers and Electronics in Agriculture*, 205: 107655.
- Fawakherji, M.; Potena, C.; Bloisi, D. D.; Imperoli, M.; Pretto, A.; and Nardi, D. 2019. UAV image based crop and weed distribution estimation on embedded GPU boards. In *Computer Analysis of Images and Patterns: CAIP 2019 International Workshops, ViMaBi and DL-UAV, Salerno, Italy, September 6, 2019, Proceedings 18*, 100–108. Springer.
- Hasan, A. M.; Soheli, F.; Diepeveen, D.; Laga, H.; and Jones, M. G. 2021. A survey of deep learning techniques for weed detection from images. *Computers and Electronics in Agriculture*, 184: 106067.
- Hu, K.; Wang, Z.; Coleman, G.; Bender, A.; Yao, T.; Zeng, S.; Song, D.; Schumann, A.; and Walsh, M. 2024. Deep learning techniques for in-crop weed recognition in large-scale grain production systems: a review. *Precision Agriculture*, 25(1): 1–29.
- Jeon, H. Y.; Tian, L. F.; and Zhu, H. 2011. Robust crop and weed segmentation under uncontrolled outdoor illumination. *Sensors*, 11(6): 6270–6283.
- Jiang, H.; Zhang, C.; Qiao, Y.; Zhang, Z.; Zhang, W.; and Song, C. 2020. CNN feature based graph convolutional network for weed and crop recognition in smart farming. *Computers and Electronics in Agriculture*, 174: 105450.
- Joiya, F. 2022. YOLO vs. Faster R-CNN for Object Detection. *International Research Journal of Modernization in Engineering Technology and Science (IRJMETS)*. Accessed: 2024-08-14.
- Keylabs. 2023. YOLOv8 vs SSD: Choosing the Right Object Detection Model. *Keylabs*. Accessed: 2024-08-14.
- Kingma, D. P.; Mohamed, S.; Jimenez Rezende, D.; and Welling, M. 2014. Semi-supervised learning with deep generative models. *Advances in neural information processing systems*, 27.
- Liu, W.; Anguelov, D.; Erhan, D.; Szegedy, C.; Reed, S. E.; Fu, C.; and Berg, A. C. 2016. SSD: Single Shot MultiBox Detector. In *European Conference on Computer Vision*, 21–37.
- Long, J.; Shelhamer, E.; and Darrell, T. 2015. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3431–3440.
- Lu, Y. 2023. CottonWeedDet12: A 12-class weed dataset of cotton production systems for benchmarking AI models for weed detection. *Zenodo*.
- Nanos, G. 2023. Object Detection: SSD Vs. YOLO. *Baeldung on Computer Science*. Accessed: 2024-08-14.
- Nasiri, A.; Omid, M.; Taheri-Garavand, A.; and Jafari, A. 2022. Deep learning-based precision agriculture through weed recognition in sugar beet fields. *Sustainable computing: Informatics and systems*, 35: 100759.
- Olsen, A.; Konovalov, D. A.; Philippa, B.; Ridd, P.; Wood, J. C.; Johns, J.; Banks, W.; Girgenti, B.; Kenny, O.; Whinney, J.; et al. 2019. DeepWeeds: A multiclass weed species image dataset for deep learning. *Scientific Reports*, 9(1): 2058.
- Parra, L.; Marin, J.; Yousfi, S.; Rincón, G.; Mauri, P. V.; and Lloret, J. 2020. Edge detection for weed recognition in lawns. *Computers and Electronics in Agriculture*, 176: 105684.
- Pirchio, M.; Fontanelli, M.; Frascioni, C.; Martelloni, L.; Raffaelli, M.; Peruzzi, A.; Caturegli, L.; Gaetani, M.; Magni, S.; Volterrani, M.; et al. 2018. Autonomous rotary mower versus ordinary reel mower—effects of cutting height and nitrogen rate on Manila grass turf quality. *HortTechnology*, 28(4): 509–515.
- Rahman, A.; Lu, Y.; and Wang, H. 2023. Performance evaluation of deep learning object detectors for weed detection for cotton. *Smart Agricultural Technology*, 3: 100126.
- Reis, D.; Kupec, J.; Hong, J.; and Daoudi, A. 2023. Real-time flying object detection with YOLOv8. *arXiv preprint arXiv:2305.09972*.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2016. Faster R-CNN: Towards real-time object detection with region proposal networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(6): 1137–1149.
- Roboflow. 2023. YOLOv8 vs. Faster R-CNN: Compared and Contrasted. Accessed: 2024-08-14.
- Sharpe, S. M.; Schumann, A. W.; Yu, J.; and Boyd, N. S. 2020. Vegetation detection and discrimination within vegetable plasticulture row-middles using a convolutional neural network. *Precision Agriculture*, 21: 264–277.
- Sportelli, M.; Pirchio, M.; Fontanelli, M.; Volterrani, M.; Frascioni, C.; Martelloni, L.; Caturegli, L.; Gaetani, M.; Grossi, N.; Magni, S.; et al. 2020. Autonomous mowers working in narrow spaces: A possible future application in agriculture? *Agronomy*, 10(4): 553.
- Steininger, D.; Trondl, A.; Croonen, G.; Simon, J.; and Widhalm, V. 2023. The cropandweed dataset: A multi-modal learning approach for efficient crop and weed manipulation.

- In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 3729–3738.
- Sudars, K.; Jasko, J.; Namatevs, I.; Ozola, L.; and Badaukis, N. 2020. Dataset of annotated food crops and weed images for robotic computer vision control. *Data in brief*, 31: 105833.
- Tanveer, A.; Chaudhry, N.; Ayub, M.; and Ahmad, R. 2003. Effect of cultural and chemical weed control methods on weed population and yield of cotton. *Pak. J. Bot.*, 35(2): 161–166.
- Tao, T.; and Wei, X. 2022. A hybrid CNN–SVM classifier for weed recognition in winter rape field. *Plant Methods*, 18(1): 29.
- United Nations. 2023. Leave no one behind. <https://unsdg.un.org/2030-agenda/universal-values/leave-no-one-behind>.
- Veeranampalayam Sivakumar, A. N.; Li, J.; Scott, S.; Psota, E.; J. Jhala, A.; Luck, J. D.; and Shi, Y. 2020. Comparison of object detection and patch-based classification deep learning models on mid-to late-season weed detection in UAV imagery. *Remote Sensing*, 12(13): 2136.
- Wang, A.; Chen, H.; Liu, L.; Chen, K.; Lin, Z.; Han, J.; and Ding, G. 2024. Yolov10: Real-time end-to-end object detection. *arXiv preprint arXiv:2405.14458*.
- Wang, C.-Y.; Bochkovskiy, A.; and Liao, H.-Y. M. 2023. 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*, 7464–7475.
- WeLASER. 2023. WeLASER Implement.
- Wu, Z.; Chen, Y.; Zhao, B.; Kang, X.; and Ding, Y. 2021. Review of weed detection methods based on computer vision. *Sensors*, 21(11): 3647.
- Xu, M.; Zhang, Z.; Hu, H.; Wang, J.; Wang, L.; Wei, F.; Bai, X.; and Liu, Z. 2021. End-to-end semi-supervised object detection with soft teacher. In *Proceedings of the IEEE/CVF international conference on computer vision*, 3060–3069.
- You, J.; Liu, W.; and Lee, J. 2020. A DNN-based semantic segmentation for detecting weed and crop. *Computers and Electronics in Agriculture*, 178: 105750.
- Yu, J.; Schumann, A. W.; Cao, Z.; Sharpe, S. M.; and Boyd, N. S. 2019. Weed detection in perennial ryegrass with deep learning convolutional neural network. *Frontiers in Plant Science*, 10: 1422.
- Zhai, X.; Oliver, A.; Kolesnikov, A.; and Beyer, L. 2019. S4l: Self-supervised semi-supervised learning. In *Proceedings of the IEEE/CVF international conference on computer vision*, 1476–1485.
- Zhang, H.; Cao, D.; Zhou, W.; and Currie, K. 2024. Laser and optical radiation weed control: a critical review. *Precision Agriculture*, 1–25.
- Zhang, H.; Zhong, J.-X.; and Zhou, W. 2023. Precision optical weed removal evaluation with laser. In *CLEO: Applications and Technology*, JW2A–145. Optica Publishing Group.
- Zhang, L.; Zhang, Z.; Wu, C.; and Sun, L. 2022. Segmentation algorithm for overlap recognition of seedling lettuce and weeds based on SVM and image blocking. *Computers and Electronics in Agriculture*, 201: 107284.
- Zhang, Z. 1996a. Developing chemical weed control and attaching importance to integrated weed management. In *Proceedings of the National Symposium of IPM in China, 1996*. China Agricultural Science and Technology Press.
- Zhang, Z. 1996b. *Weeds and their control in cotton fields*, volume 2, 1345–1349. Beijing: China Agriculture Press.