FCU ePaper

逢甲大學學生報告 ePaper

精準農業之整合植物健康診斷與灌溉系統

Precision Agriculture:

Integrating Plant Health Diagnostics and Irrigation System

作者：黃柏憲

系級：電子四甲

學號：D0987854

開課老師：梁詩婷

課程名稱：智能影像辨識系統設計實務

開課系所：電子工程學系

開課學年： 112 學年度 第 2 學期

中文摘要

這篇論文針對農業中技術勞動力短缺的問題，提出了一個創新的 AI 驅動自動化系統。該系統將先進的視覺目標識別技術與四軸機械框架整合，提供基於圖像的植物護理評估的端對端解決方案。這包括了肥料施用、自動灌溉、產品成熟度識別（以便銷售）、以及成長週期追蹤和未來發展分析等關鍵任務的自動化。系統利用約 926 張不同狀態的 Catharanthus roseus 植物圖像數據集，設計了圖像分類、目標檢測和分割的自動識別過程，專注於葉片狀況和花朵計數。綜合實驗顯示，系統的高效能表現為：葉片分割的平均準確率約為 95%，葉片枯萎分析的整體準確率為 96.20%，葉片下垂分類的準確率為 97.59%，花朵計數任務的 F1 分數為 85.37%。這一創新旨在革新作物管理實踐，提高操作效率，並通過解決勞動力短缺問題和提升植物護理效率，促進農業生產力的提升及「農業 4.0」的出現。

關鍵字：分割、四軸機械系統、目標檢測、深度學習、圖像分類、農業

Abstract

Addressing the significant skilled labor shortage in agriculture, this paper proposes an innovative AI-driven automation system that integrates advanced visual object recognition with a four-axis mechanical framework to provide an end-to-end solution for image-based plant care assessment. This includes the automation of essential tasks such as fertilization, watering, identification of product maturity for sale, and growth cycle tracking and analysis for future development. Utilizing a dataset of approximately 926 images of Catharanthus roseus plants in various states, the system designs automated recognition processes for image classification, object detection, and segmentation, focusing on leaf conditions and flower counting. Comprehensive experiments demonstrate the system's high efficacy, with an average accuracy of approximately 95% for leaf segmentation, 96.20% overall accuracy in leaf withering analysis, 97.59% accuracy in leaf drooping classification, and an F1-score of 85.37% for the flower counting task. This innovation aims to revolutionize crop management practices, enhance operational efficiency, and contribute to the advancement of agriculture by addressing labor shortages and improving plant care efficiency, ultimately promoting agricultural productivity and the emergence of Agriculture 4.0.

Keyword：Agricultural, Deep Learning, Segmentation, Object Detection, Image Classification, Four-Axis Mechanical System

目 次

1. Introduction

The agricultural sector faces a significant challenge marked by a scarcity of skilled labor, prompting the exploration of innovative solutions to sustain and enhance productivity. In response to this pressing issue, this study introduces an advanced AI-driven automation system specifically tailored for agricultural applications. This system aims to revolutionize traditional crop management practices by leveraging cutting-edge technologies, particularly in the domain of flower cultivation. Figure 1 illustrates the demographic composition of the agricultural workforce, revealing a pronounced trend towards an aging population [1], with 63% of the laborers being above 50 years old. This underscores the pressing need for modernization and automation to mitigate labor shortages.

AI Pal [2], a Taiwanese startup, aims to enhance agricultural productivity by leveraging drones and AI. Their cloud-based platform provides services including drone surveying, image analysis, and agricultural management advice. This technology swiftly identifies field stressors,

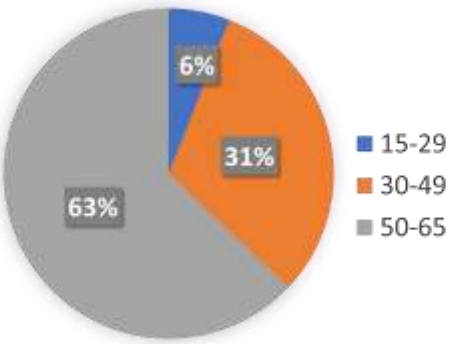

Figure 1: Agricultural population age distribution pie chart

diseases, and pests, facilitating informed decision-making on irrigation, fertilization, and pest control. AI Pal's partnerships with local governments in Taiwan have yielded significant improvements in various agricultural fields, such as rice paddies, pineapple orchards, and mango or- chards, aiming to reduce crop losses and increase efficiency. By pioneering smart agriculture technologies, AI Pal has successfully contributed to the transformation and modernization of Taiwan's agricultural sector.

Inspired by recent technological advancements, this paper proposes utilizing deep learning models to introduce an automation system for real-time plant care assessments, advancing towards Agricultural 4.0. The system is de-

signed to identify common plant defects, such as yellow- ing leaves, discrepancies in flower quantity, and instances of leaf curling or drooping, using visual-based identifica- tion techniques. These techniques are integrated into afour-axis mechanical slide rail system that enables auto- matic watering and fertilization for precision agriculture.In brief, this hardware system is controlled by an Arduino Uno and a CNC shield, enabling bidirectional communi- cation between hardware devices and software automation for accurate plant health diagnostics and efficient irriga- tion.

While the proposed system initially targets flower cultivation, its potential extends beyond this domain, laying the groundwork for future research in expanding recognition systems to accommodate various plant species and agricultural practices. By offering a scalable and sustainable solution to labor shortages in agriculture, this work aims to significantly enhance operational efficiency while mitigating labor constraints, fostering a paradigm shift towards increased agricultural productivity.

In short, the main findings of our study are highlighted as follows:

1. Compilation of a varied dataset, consisting 926 images of Catharanthus roseus plants.
2. Introduction of a comprehensive end-to-end recognition process, incorporating image classification, object detection, segmentation algorithms, and quantification of flowers.
3. Development and utilization of a four-axis mechanical slide rail system for capturing images, along with watering and fertilization actions.
4. Seamless integration of automated recognition technologies to activate the hardware components promptly.
5. Extensive experimental work and thorough analysis conducted to assess the performance and effectiveness of the developed system.

In the subsequent sections, the paper provides a comprehensive overview of the research, presenting a thorough review of relevant literature in Section 2. Following this, the proposed methodology is elaborated in Section 3, and the experimental setup is detailed in Section 4. Key findings and empirical evidence are discussed in Section 5. Through this structured approach, the paper aims to offer a cohesive understanding of the study and its implications within the broader context of agricultural innovation, concluding in Section 6.

2. Literature Review

Despite significant advancements in image processing techniques over the past few decades, the development of automated plant monitoring and analysis research remains constrained, primarily due to the lack of a dataset for ex- perimental evaluation. Section 2.1 offers a concise review of previous studies dedicated to automated plant analysis, emphasizing their major research outcomes. Furthermore,a research gap is identified in Section 2.2, underscoring the necessity for the proposed system.

2.1. Automated plant recognition system

Precision agriculture has been embraced by research groups worldwide, aiming to leverage cutting-edge tech- nologies to enhance the accuracy and efficiency of agricul- tural practices. This section highlights and discusses the applications and innovations of significant works, underscoring the importance of precision agriculture technolo- gies in promoting a more informed and precise approach to managing plant health and water usage. Some related works published previously are reviewed in this section,and a brief summary of these papers is provided in Ta-ble 1.

A recent study conducted by Toğaçar et al. by [3], whereby they focus on the classification of flower species, employing a hybrid approach that combines feature selec-tion techniques with CNN models. This work capitalizes on an openly accessible dataset that contains a substantial collection of 4242 flower images. These images were parti- tioned into two parts, where the training and test sets are with the data distribution of 80% and 20%, respectively. To discern the most influential features for classification, the study incorporated feature selection techniques such as the genetic algorithm (GA) [8] and the tabu search al- gorithm [9] . Subsequently, a Support Vector Machine (SVM) classifier [10] was deployed to categorize the flowerimages into five distinct classes. The outcome was nothing short of remarkable, boasting an impressive overall accu-racy rate of 98.91%. This was achieved by meticulously selecting the top 2500 features to ensure optimal feature selection. Nonetheless, there appears to be a lack of suf- ficient detail regarding the plant, particularly concerning aspects such as stem structure or leaf characteristics. The focus primarily revolves around the identification of flowerspecies.

On a related note, [4] proposes an innovative hybrid method centered on CNN architecture for precise clas- sification of flower species. This strategy harnesses the capabilities of two pre-trained CNN models, specifically VGG-16 [11] and AlexNet [12], originally trained on the ImageNet dataset [13], to streamline the process of feature extraction and selection, resulting in improved classifica- tion accuracy. The research utilizes two distinct datasets: Flower17, comprising 1360 images representing 17 diverseflower species, and Flower102, which encompasses a more extensive collection of 8189 images spanning 102 distinct

flower categories. Notably, the study reveals that by opti- mizing feature selection using specific feature layers from AlexNet for Flower17 and selective feature layers from both AlexNet and VGG-16 for Flower102, classification ac- curacy achieves impressive results of 96.39% and 82.68%,

Table 1: Summary of the existing methods proposed for the automated recognition plant monitoring task

No.	Ref.	Task	Target	Images	Labels	Model	Accuracy
1	[3]	Classification	Flower	4242	5	Multiple	98.91%
2	[4]	Classification	Flower	9549	119	AlexNet & VGG-16	96.39%
3	[5]	Object Detection	Flower & Fruit	2415	3	FaceNet	97.70%
4	[6]	Semantic Segmentation	Leaf & Stem	5460	6	PSegNet	94.52%
5	[7]	Classification	Crop species & Diseases	54,306	38	AlexNet & GoogLeNet	99.35%

respectively, highlighting the effectiveness of the proposed methodology. However, since this paper primarily concentrates on the technical intricacies of feature selection and classification, it lacks real-world implications to evaluate the effectiveness of the proposed method.

On the other hand, [7] utilized CNN model was effectively employed to accurately identify a wide range of 26 distinct crop diseases solely from individual leaf images. This innovative solution holds substantial poten- tial for advancing crop management practices and ensuring global food security. To achieve this, the study harnessed a publicly available dataset known as the PlantVillage dataset [14], comprising an extensive collection of 54,306 images portraying both diseased and healthy plant leaves. The CNN was meticulously trained using pre-trained CNN models, specifically AlexNet [12] and GoogLeNet [15], enabling it not only to classify diverse crop species but also to accurately determine disease statuses across a wide spectrum of 38 categories. These encompassed 14 crop species and 26 specific diseases. Remarkably, the model demonstrated exceptional performance, achieving an impressive accuracy rate exceeding 99%. This accomplishment underscores the significant promise of deep learning in addressing critical agricultural challenges. However, it is highlighted that there is a significant drop in accuracy to 31.4% when tested on real-world images compared to the impressive 99.35% reported. This emphasizes the need for a more diverse training dataset to improve generalization accuracy. Additionally, the model's limitation to classifying single upward-facing plant leaves restricts its applicability, as diseases may manifest on different plant parts or from varied angles.

In a higher-dimensional context, the study in [6] delved into 3D point cloud segmentation, introducing an innovative approach called PSegNet, enabling simultaneous semantic and instance segmentation in the 3D space. This novel method capitalizes on a 3D dataset acquired through laser scanning, which was initially established by Conn et al. [16, 17]. Notably, the dataset encompasses three dif- ferent crops—tomato, tobacco, and sorghum—capturing their growth status in various environmental conditions. Impressively, the research yielded commendable results in semantic segmentation, with quantitative performance metrics exceeding 85% for the entire test dataset. Remarkably, the segmentation results exhibited a higher level of accuracy for leaves compared to stems. This study represents a significant advancement in enhancing the precision of plant point cloud segmentation, holding promising implications for diverse applications in agriculture and environmental sciences.

A more recent study conducted in [5] introduces an innovative approach for detecting fruits and flowers in a strawberry field, employing the Faster R-CNN architecture [18] with backbone model of ResNet-50 [19]. Specifically, they harness the FaceNet model [20] to learn feature embeddings, enabling the clustering of strawberry objects. This study employs two distinct clustering techniques, namely Density-Based Spatial Clustering of Applications with Noise (DBSCAN) [21] and threshold clustering, to categorize strawberry objects into three groups: flowers, unripe fruits, and ripe fruits. Remarkably, the overall accuracy achieved with DBSCAN is 99.26%, while threshold clustering yields an accuracy of 97.70%, demonstrating the efficiency of both clustering methods in accurately detecting fruits and flowers in the strawberry field. While the study relies solely on an object detection system, it may struggle to accurately identify diseases or defects within plants. Object detection systems are limited in their ability to precisely delineate subtle abnormalities. In contrast, the semantic segmentation method can offer a more detailed classification up to the pixel level, enabling finer-grained identification of affected regions within plant images.

2.2. Research Gap

Despite the high accuracy of plant classification or detection demonstrated by the studies mentioned above, automation in plant monitoring and management remains an area ripe for further research and development. Although existing research has shown promising results in identifying common plant defects, crop species, and disease determination, there are still significant gaps in the implementation of automated watering and fertilization systems. Integrating these systems with plant health diagnostics is essential for achieving precision agriculture. Furthermore, there is a clear need for ongoing research efforts aimed at refining automated plant care systems to enhance their feasibility and applicability in real-world agricultural settings.

3. Proposed Method

The main objective of this work is to develop a sustainable plant monitoring and irrigation system that assesses growth conditions to determine their suitability for sale. Initially, a software algorithm is implemented to monitor if the leaves of the plants exhibit a yellowish color, indicating the need to increase the frequency of watering and apply fertilizers. In cases where leaves show signsof drooping, this suggests either a lack of water or ex- cessive exposure to sunlight, necessitating adjustments in watering quantity or advising farmers to reduce the du- ration of plant exposure to sunlight. Consequently, hard- ware devices are constructed for watering and fertilizing purposes. When the number of flowers reaches a prede- termined threshold, farmers are notified that the plant is ready for sale. Figure 2 depicts the overview flowchart of the proposed pipeline, which consists of seven main steps:
(a) machine platform construction, (b) data collection, (c) leaf detection, (d) machine control, (e) flower detection, (f) determination of plant readiness for sale, (g) interactive platform design and data analysis. The following subsections will elucidate each step with detailed procedures.

3.1. Machine Platform Construction

The hardware machine utilized is a customized four-axis mechanical slider with dimensions of 48 cm×46 cm×44 cm. To control the motion of the rail system precisely, three stepper motors are used to navigate movement along the XYZ axes. A lightweight camera is mounted on the rotatable Z-axis at a tilted angle (45°) facing downward, facilitating 360° capture of plant images. At the bottomof the Z-axis, a stepper motor is installed to manage the rotation. For watering or fertilizing purposes, a spray noz- zle is positioned on the Z-axis. Considering the impact of weight and torque, the use of a high-torque stepper motor is considered essential. This four-axis mechanical slider system is governed by an Arduino Uno paired with a CNC Shield. MATLAB software acts as the image recognition tool, ensuring seamless bi-directional communication with the Arduino Uno. Additionally, the designed machine platform encompasses functionalities such as transmitting coordinate information, sending arrival signals, performing image recognition, and managing the spray nozzle switch, thus achieving a fully automated control system for plant growth assessment and care. Figure 3 presents an exploded view of the constructed machine for enhanced visualiza- tion.

3.2. Data Collection

All images evaluated in this experiment were self-collected data. To capture a wide range of images for each plant, two cameras were used: one mounted on the rotating Z-axis of the hardware mechanism, strategically positioned at a 45°, and another tilted 90° above the plant head, capturing photos of the plant from different orientations by rotating on the axis (a). Two Logitech C922 Pro Stream webcams with detailed specifications are tabulated in Table 2. The experimental environment is enclosed with a white background, and constant lighting conditions are maintained using ILL-DP2454-4CR in conjunction with LED light panels.

Table 2: Specification and configuration of the camera to collect the experimental data

Feature	Description
Model	Logitech C922 Pro Stream
Resolution	1920×1080
Frame/ second	30fps
Camera mega pixel	3
Focus type	Autofocus
Diagonal field of view	78°
Digital zoom	1.2x

To ensure a sufficient amount of experimental data could be collected, Catharanthus roseus was chosen for its ease of growth and low maintenance requirements. With minimal care, such as consistently moist soil and balanced fertilization, this plant can produce showy and vibrant red flowers. In total, 15 plants were utilized to address diverse variables influencing plant growth and encompass different growth stages, resulting in 926 images for the experimental data. Notably, the captured features include discoloration, curling, and drooping of leaves, the presence or absence of flowers, and the plant's transition from seedling to flowering.

3.3. Leaf Detection

After acquiring the image data, the leaf regions in the images will be identified for analysis of drooping and withering conditions. To achieve this, two specialized deep learning techniques are utilized: semantic segmentation is applied to identify withering conditions, while object classification is employed to detect drooping conditions. A more detailed explanation of the methods, along with their respective intuition and insights, is provided below.

- Withering determination:
  In brief, determining plant withering conditions involves two main steps: leaf localization and color pixel thresholding. For leaf localization, the deep learning technique known as semantic segmentation, specifically the DeepLabV3+ [22] algorithm with the ResNet-18 [19] CNN architecture, will be used. DeepLabV3+ employs an encoder-decoder architecture, where the encoder handles information at different grid scales through dilated convolutions, and the decoder focuses on refining segmentation results to highlight object boundaries. The inclusion of Atrous Spatial Pyramid Pooling (ASPP) allows the algorithm to gather multi-scale contextual information while preserving spatial accuracy, facilitating precise pixel-level classification for detailed leaf geometry analysis.

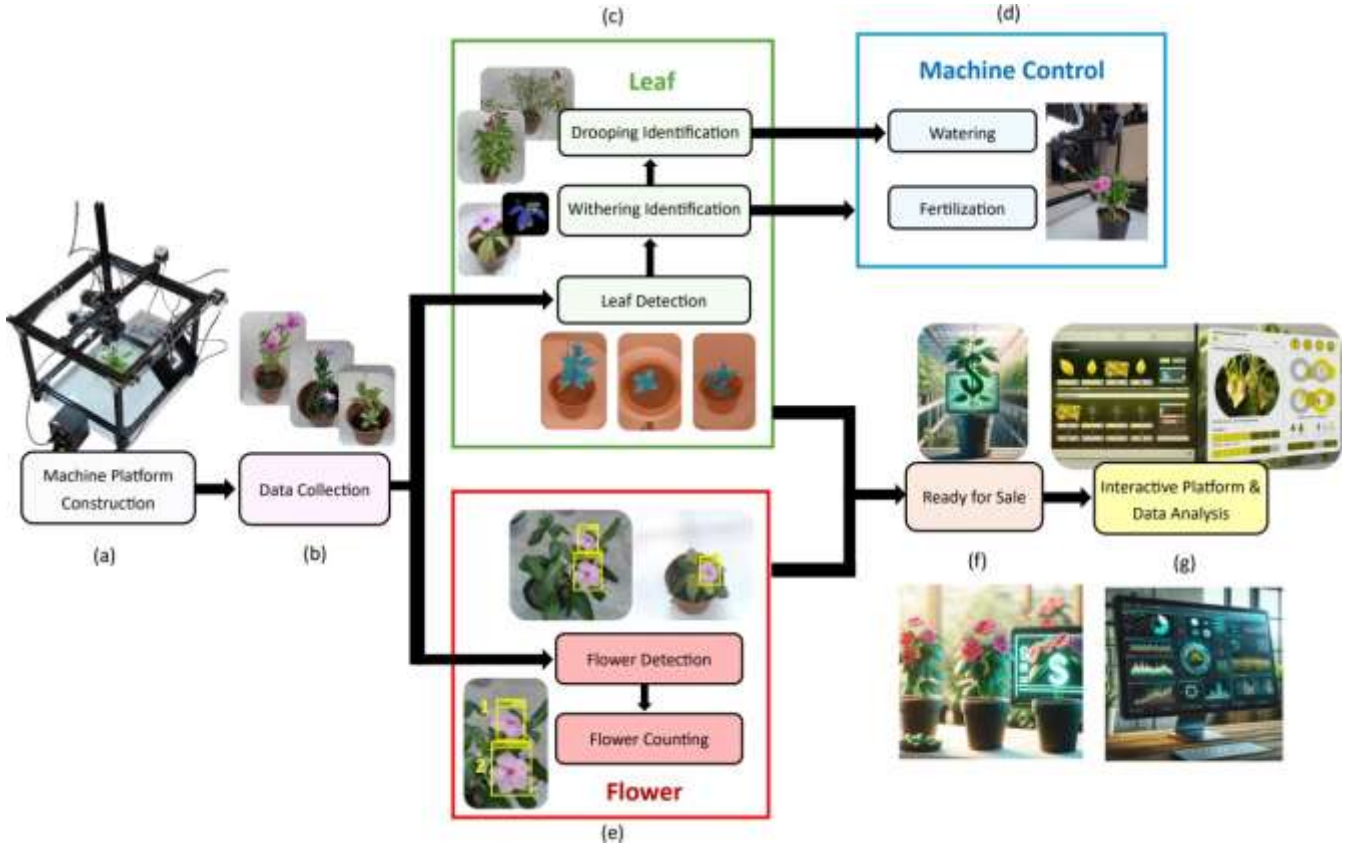

Figure 2: The proposed plant monitoring and irrigation system encompasses seven main steps: (a) machine platform construction, (b) data collection, (c) leaf detection, (d) machine control, (e) flower detection, (f) determination of plant readiness for sale, and (g) interactive platform design and data analysis.

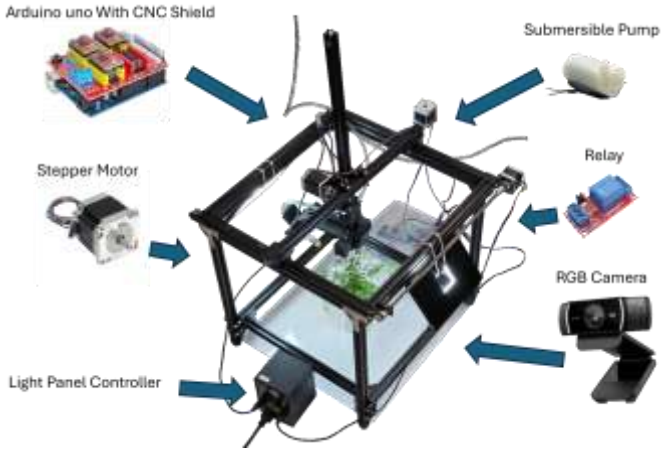

Figure 3: The exploded view of the hardware mechanism system

The process begins with pixel annotation of the images, categorizing them into two groups: background and leaf. This task is accomplished using the Image Labeler application in MATLAB 2022b software, with an annotation example shown in Figure 4 (a), where leaf pixels are marked in blue and background pixels in orange. After training the DeepLabV3+ model for leaf segmentation, the model efficiently identifies target regions for the next stage, which involves color thresholding to determine withering.

The following analysis focuses solely on the yellow channel of the pixels. Thresholds for the red, green, and blue channels are set to above 130 for red and green, and below 145 for blue, under uniform lighting conditions, to single out yellow pixels. Subsequently, the proportion of yellow pixels is calculated to assess withering. A plant is considered to be in a withering condition if the yellow proportion exceeds 20% of the total image area; it is deemed healthy if the proportion is below 20%. Figure 4 (b) offers a visual representation of the results following the application of leaf segmentation and color thresholding techniques. The purple regions indicate the identified yellow segments.

- Drooping determination:
  To identify curled or drooping leaves in plants, an image classification approach employing transfer learning is implemented. The pre-trained ResNet-18 model is selected as the CNN architecture, and the model is fine-tuned with images illustrating curled and non-

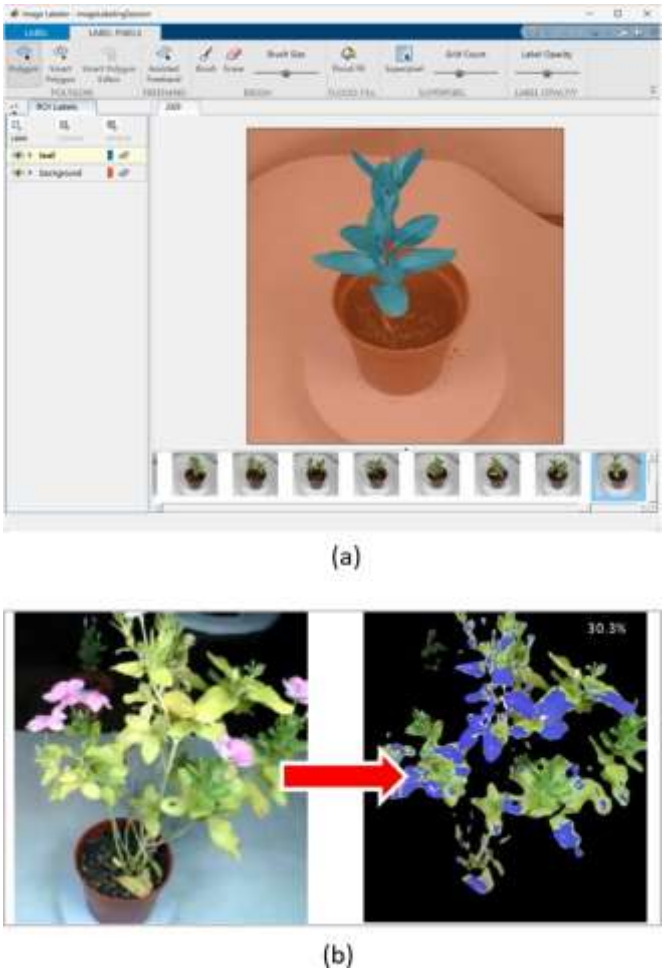

Figure 4: (a) Leaf annotation in the Image Labeler application, where leaf pixels are marked in blue and background pixels in or-ange, and (b) the process of executing leaf segmentation and pixel thresholding to identify withering conditions.

curled conditions. It is important to note that the pre-trained model was previously trained on Ima- geNet [13], which consists of more than 1 million im- ages across 1000 categories. The assessment of droop-ing relies on images captured from various angles,leading to the proposal of a simple voting mecha-nism to determine whether a plant is in a drooping or healthy condition. An image displaying both healthy and drooping conditions is shown in Figure 5.

3.4. Machine Control

To construct an automated plant health monitoring system, the designed machine incorporates functions for both watering and fertilizing. Specifically, the system initiates watering when the yellowish discoloration of the leaf region, indicative of withering, exceeds 20% of the entire image. Conversely, fertilization is activated solely when the yellowish discoloration in the leaf region surpasses 40%. The criteria of triggering the watering and fertilizing actions can be summarized as:

Figure 5: The exemplar of plants with drooping and healthy conditions.

$$\text{Withering} = \begin{cases} 20\% < x < 40\% & \text{, watering;} \\ x \geq 40\% & \text{, fertilizing and watering} \end{cases} \quad (1)$$

Additionally, the presence of a drooping condition also triggers the watering function. That is to say that when the leaves presented mild yellowing and curling, it is a sig- nal of deficiency in nutrients and water, thus necessitating fertilization and watering.

In the context of the irrigation system, the nozzles designated for watering and fertilizing are precisely positioned at the z-axis manipulator rod angled at 45°, enhancing the efficiency of the rotating spray. The system utilizes an innovative design incorporating a three-way pipe equipped with a check valve, allowing for the installation of just one nozzle at the end. This setup effectively blocks the reverse flow of water and liquid fertilizer. Positioned at this crucial juncture are two 5V submersible motors, connected to a high-voltage trigger relay that facilitates the switching of motor functions. MATLAB software is instrumental in issuing command codes to the Arduino Uno, triggering the system's activation. The mechanisms for watering and fertilizing are depicted in Figure 6.

3.5. Flower Detection

To enable accurate flower counting, an object detection method is utilized to initially identify flower positions, followed by counting the detected flowers. The foundation for

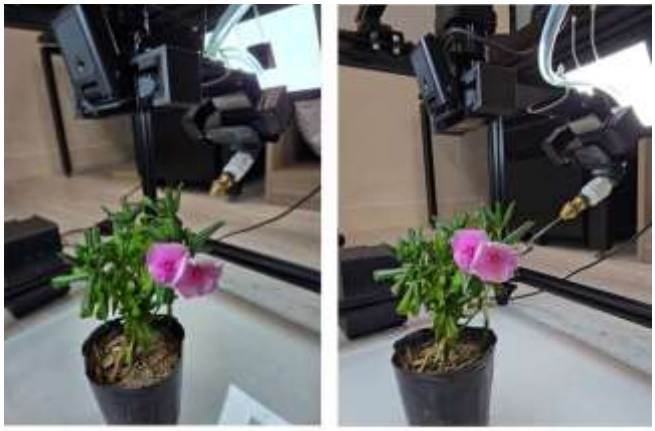

Figure 6: (a) The Z-axis manipulator rod, equipped with a nozzle for both watering and fertilizing functions, (b) triggering of the watering action, and (c) activation of the fertilizing action.

the flower detection algorithm is the Image Labeler tool in MATLAB software, as demonstrated in Figure 7 (a). It is crucial to acknowledge that the target objects are annotated within rectangular bounding boxes, contrasting with the pixel-based annotations used in Section 3.3 for leaf segmentation.

The Faster R-CNN model [23], with ResNet-50 [19] as the backbone architecture, is deployed for object detec-tion. This algorithm introduces a Region Proposal Net- work (RPN), significantly boosting the model's capacity to generate candidate regions. Additionally, Faster R-CNN adopts an end-to-end training methodology, streamlining and enhancing the training process for better efficiencyand integration. The execution of flower detection, inclu- sive of the confidence scores for each detection, is shownin Figure 7 (b).

Subsequently, a straightforward count of the detected bounding boxes is conducted to determine the number of flowers present on each plant. To aggregate the number of flowers, an average of the flowers identified is calculated for each image taken of the plant. It is important to note that images in which no flowers were detected are excluded from the calculation to avoid potential inaccuracies stemming from angles where flowers are not visible.

3.6. Plant Readiness for Sale

To ascertain whether a plant is ready for sale, it must meet three fundamental criteria: it should not be drooping or withering, and it must have a minimum of five flowers. The evaluation of these criteria involves a detailed analysis of the plant's physical condition and its flowering state.

1. Non-drooping: The plant must exhibit a healthy pos- ture without any signs of drooping leaves or stems. Drooping is often an indicator of insufficient water uptake or overexposure to heat, which can affect the plant's overall health and aesthetic appeal.

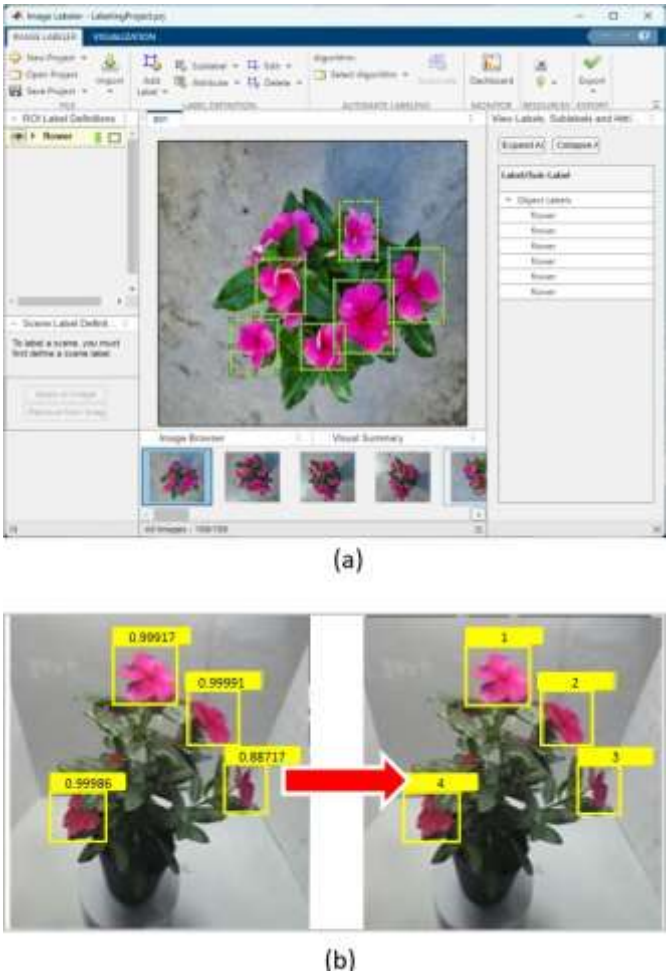

Figure 7: (a) Annotation of flower bounding boxes in the Image Labeler application, and (b) the process of conducting flower detection and counting.

2. Non-withering: For a plant to be considered non- withering, its overall appearance must be vibrant and free from signs of decay or dehydration. The spe- cific threshold for withering determination in this con- text is set at a yellowing percentage of less than 20%. This means that less than 20% of the plant's foliage should show yellowing, which is commonly associated with nutrient deficiencies, over-watering, or under- watering.

3. Minimum flower count: A plant must have at least five fully developed flowers to qualify for sale. This criterion ensures that the plant is attractive to poten- tial buyers and indicates a healthy growth stage. The flower count serves as a direct indicator of the plant's reproductive success and overall vitality.

3.7. Interactive Platform and Data Analysis

Once the models for evaluating leaf condition and flower count are fully trained, they are deployed to assess the plant's growth status in real time. During this assessment

phase, a total of 11 images are captured for each plant (10 at 45° angles and 1 at a 90° angle) to ensure a comprehensive analysis. The condition of the leaves, including any signs of withering and drooping, is then aggregated and displayed on the GUI, indicating whether the plant requires fertilization and watering interventions. Simultaneously, the GUI also presents the total flower count, providing a holistic view of the plant's health and flowering status on the same screen. An example of the designed GUI is shown in Figure 8. Specifically, the left part of the window lists the criteria for judging withering, water-ing, and selling conditions, while the right part shows theon-the-spot leaf segmentation, flower detection, and flower counting results.

Upon completing these recognition tasks, the GUI transitions to a final screen that provides an in-depth analysis of the plant's overall growth condition, offering guidance on whether the plant is sufficiently healthy and well-developed to be considered ready for sale. This final interface is designed to make interaction with the system as intuitive as possible, ensuring that users are thoroughly informed about the health and marketability of the plant. By offering a detailed yet accessible overview of the plant's condition, the enhanced GUI supports informed decision-making about its suitability for sale. In short, this interactive GUI aims to provide a seamless and enlightening user experience, effectively bridging the gap between the complexities of recognition technology and user-friendly presentation.

4. Experiment Setup

This section elaborates on the details of the dataset acquired and the experimental configuration used for train- ing the models for various tasks (i.e., leaf segmenta- tion, withering identification, drooping identification, and flower counting). Additionally, the performance metrics adopted for each task are elucidated in this section.

4.1. Database Description

The database used in this experiment comprises 926 images collected from 15 varieties of chrysanthemum plants. These images represent various growth conditions of the plants, as detailed in Table 3. Specifically, the dataset includes 49 images of withering plants, 12 images of withering plants that contain flowers, 404 images of plants that contain flowers, 107 images of drooping plants, 69 images of drooping plants that contain flowers, and 285 images of seedlings with leaves only. To provide a clearer visualization of the images used in the experiments, some sample images are shown in Figure 9. Specifically, the figure portrays the three different health conditions captured at both 45°

Table 3: The experimental plant samples with different health conditions

Condition	Number of samples
Withering	49
Withering and contain flowers	12
Contain flowers	404
Drooping	107
Drooping and contain flowers	69
Seedlings (leaves only)	285
Total	926

4.2. Experiment Configuration

The experiments were conducted on a platform equipped with an Intel(R) Core(TM) i7-12700H 2.3 GHz CPU and an NVIDIA® GeForce RTX™ 3060 Laptop GPU, with 16GB of memory. MATLAB 2022b was used for training the image recognition models. The input size for each CNN model varied slightly depending on the se- lected backbone architecture, necessitating resizing of the database. The detailed training configurations for each task are provided in Table 4, which includes information on the tasks, approaches, backbone architectures, input sizes, learning rates, maximum epochs, and mini-batch sizes. In short, the three tasks were trained based on ResNet backbone architectures with input sizes of 400×400×3 or 640×480×3.

4.3. Performance Metric

Since the developed system involves three tasks—leaf segmentation, drooping determination, and flower counting—various metrics were used to evaluate the performance of the algorithm for each task. The rationale behind each metric, along with the corresponding mathematical equations, is detailed in the following subsections.

4.3.1. Leaf segmentation

To evaluate the performance of leaf segmentation, four metrics were adopted: mean accuracy (*mAcc*), mean Intersection over Union (*mIoU*), weighted Intersection over Union (*WIoU*), and mean Boundary F1-Score (*mBF*).

1. *mAcc*:
   This metric assesses the accuracy of correctly identifying pixels for each class (e.g., background and leaves). It provides an overall measure of how well the model is performing in terms of pixel-wise classification and can be expressed as:

$$mAcc = \frac{1}{S}\sum_{s=0}^{S}\left(\frac{\sum_{c_i} TP_{s,c_i}}{\sum_{s,c_i} TP_{s,c_i} + FP_{s,c_i}}\right)_i \quad (2)$$

and 90° angles.

where S represents the total number of images involved in the testing evaluation and C denotes the total number of classes (i.e., 2). On the other hand,

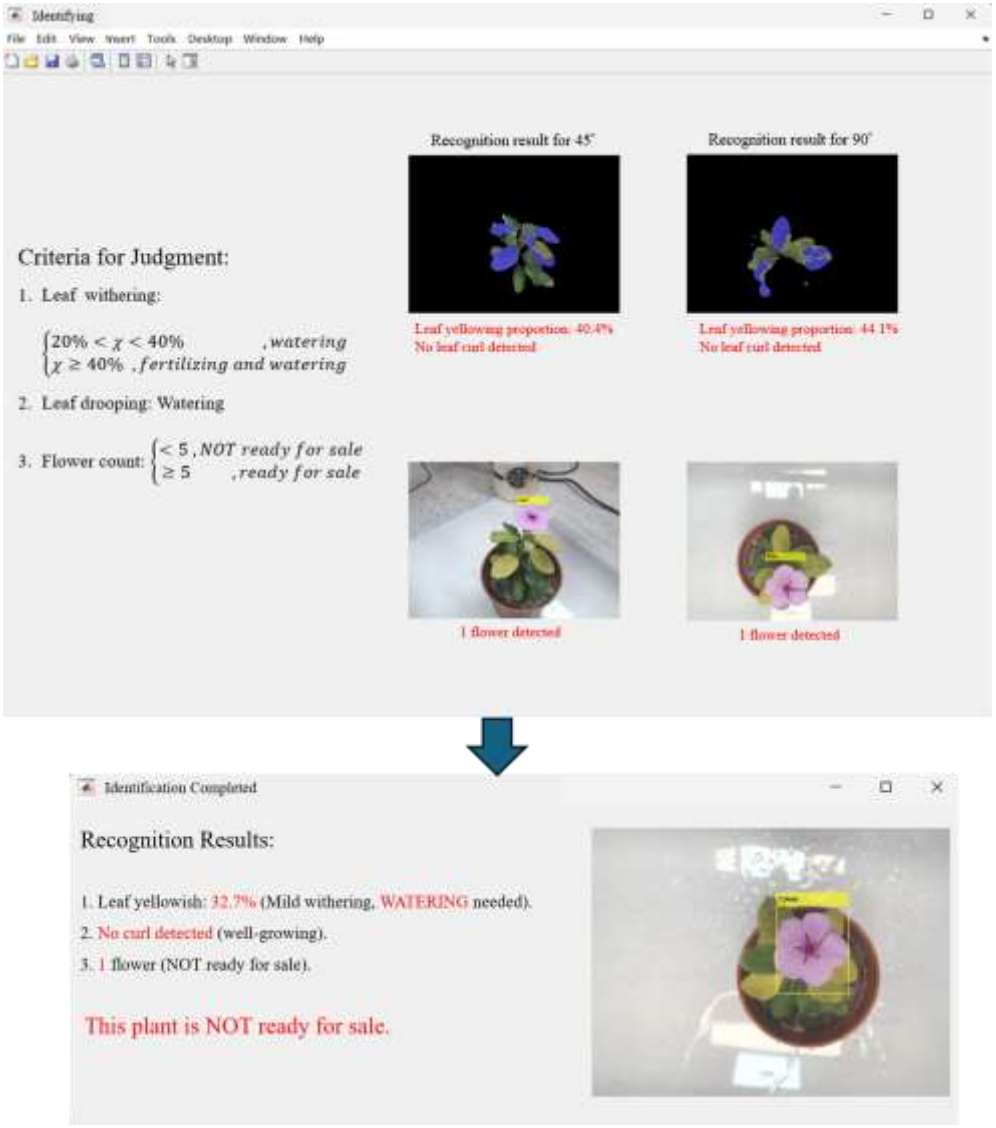

Figure 8: The interactive GUI displays the results of leaf segmentation, flower detection, and flower counting, providing intuitive insights for further actions such as watering, fertilizing, or determining readiness for sale.

TP_{s,c_i} and FP_{s,c_i} are the number of true positive and false positive instances for a specific class c_i within the image s, respectively.

2. *mIoU*:
IoU measures the degree of overlap between detected leaves and the ground truth annotation regions. It calculates the ratio of the intersection area of the detected object region and the ground truth annotation region to their union area. *mIoU* averages this value across both the background and leaves classes, providing a comprehensive measure of segmentation accuracy.

$$\frac{1}{} \sum \quad \frac{1}{} \sum \frac{TP_{s,c_i}}{}$$

where FN_{s,c_i} is the number of false negative instances for a given class c_i within the sample s.

3. *WIoU*:
This metric is the weighted average of *IoU* for each class, where the *IoU* for each class is weighted by its pixel count(P_{s,c_i}). This ensures that classes with more pixels contribute more to the overall *IoU*, pro- viding a balanced evaluation that accounts for class imbalances.

$$WIoU = \frac{1}{S} \sum_{s=0}^{S} \sum_{c_i=0}^{C} \frac{P_{s,c_i} \times TP_{s,c_i}}{\sum_{c_i=0}^{C} P_{s,c_i})(TP_{s,c_i} + FP_{s,c_i} + FN_{c_i})} \quad (4)$$

$$mIoU = \sum_{s=0}^{S} \sum_{c_i=0}^{|C|} TP_{s,c_i} + FP_{s,c_i} + FN_{s,c_i} \quad (3)$$

4. mBF:
This metric evaluates the boundary quality of the segmentation, measuring how well the model delineates

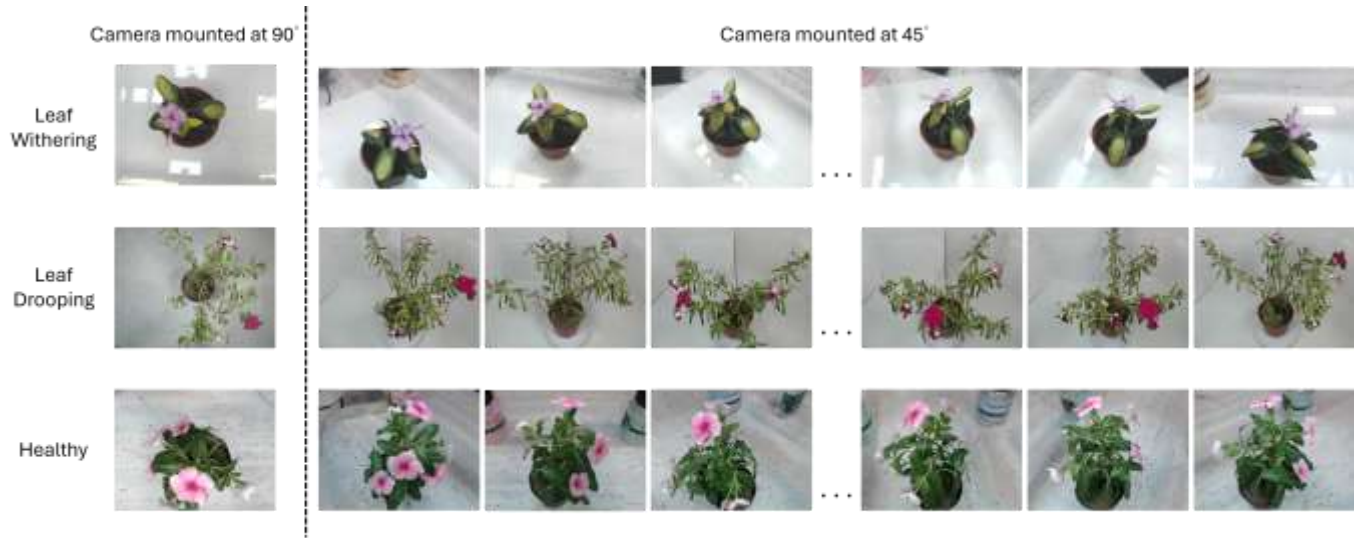

Figure 9: Sample images with three different healthy conditions captured at two angles

Table 4: The training configuration for the three tasks (i.e., leaf segmentation, drooping determination, and flower detection) involved in the plant health diagnostics and irrigation process

	Task	Approach	Backbone Architecture	Input Size	Learning Rate	Max Epoch	Mini-batch Size
1	Leaf segmentation	DeepLabV3+	ResNet-18	400×400×3	0.0001	5	2
2	Drooping determination	Binary classification	ResNet-18	640×480×3	0.0001	8	2
3	Flower counting	Faster R-CNN	ResNet-50	400×400×3	0.00001	110	1

the boundaries of the segmented regions. It is particularly useful for tasks where precise boundary detection is crucial.

$$mBF = \frac{1}{S}\sum_{s=0}^{} \left(\sum_{s,c_i} \frac{2TP_{s,c_i}}{2TP_{s,c_i} + FP_{s,c_i} + FN_{s,c_i}} \right) \quad (5)$$

4.3.2. Drooping and Withering Determination

To determine the drooping condition, this study em- ploys a binary classification approach. After identifying leaf regions, they are categorized into two groups: non- drooping and drooping. For yellowing analysis, a three- category classification method is utilized, categorizing yellowing severity into three levels: no yellowing, mild yellowing, and severe yellowing. When assessing the performance of the classification models, four key metrics are employed: accuracy, F1-score, precision, and recall.

1. Accuracy:
   It measures the ratio of correctly classified samples to the total number of samples. It provides an overall effectiveness of the model across all classes (binary or three-class).

where c_{ii} represents the number of true positives for class i, while c_{ij} denotes the number of samples incorrectly classified as class i instead of class j. k is the total number of classes.

2. Precision:
   It reflects the model's ability to accurately predict positive instances. In the context of binary classification (drooping/ non-drooping), it measures how many of the identified drooping instances are actually drooping. For three-class classification (yellowing), it indicates how accurately each class (no yellowing, mild yellowing, severe yellowing) is identified.

$$\text{Precision} = \sum_{i=1}^{k} \frac{c_{ii}}{\sum_{j=1}^{k} c_{ji}} \quad (7)$$

3. Recall:

It measures the model's ability to identify all actual positive instances. In binary classification, it evalu- ates how many of the actual drooping instances are correctly identified. For three-class classification, it assesses how well each class is detected.

$$\text{Accuracy} = \frac{\sum_{i=1}^{k} C_{ii}}{\sum_{i=1}^{k}\sum_{j=1}^{k} C_{ij}} \quad (6)$$

$$\text{Recall} = \sum_{i=1} \frac{C_{ii}}{\sum_{j=1}^{k} C_{ij}} \quad (8)$$

$\sum \quad \sum$ $\sum$

4. F1-score:
   It is the harmonic mean of precision and recall, offering a balanced assessment of the model's performance. It is particularly useful in situations where the class distribution is imbalanced, as it considers both false positives and false negatives.

$$\text{F1-score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \quad (9)$$

4.3.3. Flower Counting

Note that this flower counting process initially relies on the object detection algorithm (i.e., Faster R-CNN), for which the commonly used metric is IoU. However, since the ultimate aim of this flower counting task is to accu- rately assess the number of flowers present in a video cap- turing the plant from a 360° perspective, resulting in a to- tal of 11 images per video, additional analysis is required after detecting the flowers. Consequently, three metrics are employed in this flower counting task: precision, re- call, and the F1-score. The metrics offer intuitive insights into quantifying the error between the number of flowers identified by the model and the actual number. The mathematical formulas for each metric are provided below:

1. Precision:
   This metric indicates the accuracy of the model in correctly identifying flowers. A high precision value suggests that when the model predicts the presence of a flower, it is correct most of the time.

$$\text{Precision} = \frac{TP}{TP + FP} \quad (10)$$

2. Recall:

   It indicates the ability of the model to capture all the flowers present in the video. A high recall value suggests that the model can identify most of the flowers present.

$$\text{Recall} = \frac{TP}{TP + FN} \quad (11)$$

3. F1-score:

   It offers a balanced assessment of the model's performance in terms of both identifying flowers accurately and capturing all the flowers present.

$$\text{F1-score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \quad (12)$$

5. Experimental Result and Discussion

To demonstrate the effectiveness of the proposed framework in monitoring plant growth conditions, the focus lies on the performance analysis of CNN model training and evaluation, as reported and discussed in Sections 5.1, 5.2,5.3, and 5.4. Additionally, a summary and the limitations of the proposed system are finalized in Section 5.5. It is important to note that this study utilizes a newly collected dataset of plant specimens, encompassing tasks such asleaf segmentation, leaf yellowing/ withering identification, leaf curling classification, and flower counting. Therefore, comparisons with existing literature may not be directly applicable to this research. The data collection involves 15 plants exhibiting various growth conditions, including seedlings and mature plants, varying degrees of leaf yellowing, presence or absence of leaf curling, and different flower bud counts. Additionally, the dataset includes plant photos captured from different angles.

5.1. Results of leaf segmentation

In this section, a dataset of 395 images was extracted from the database for leaf segmentation, and it was split into training and testing sets in an 8:2 ratio for cross-validation. This process was repeated five times to ensure thorough testing of all data. The results of leaf segmentation are presented in Table 5, indicating satisfactory performance across all validation metrics: $mAcc$ at 95.01%, $mIoU$ at 0.9183, $WIoU$ at 0.9747, and mBF at 0.9373. These results indicate the effectiveness of the proposed framework in leaf detection.

However, it is important to note that leaf pixels oc- cupy less than one-fifth of the entire image area. There- fore, when calculating $WIoU$, the background IoU score is weighted higher, while $mIoU$ provides a fairer evaluation by emphasizing leaf scores, thereby mitigating the issue of inflated IoU due to the dominance of background pixels in the image. Consequently, based on the $mIoU$ reportedin Table 5, the proposed method typically outperforms in leaf segmentation tasks, with an average score of 0.9183, peaking at 0.9434 during the first validation. However, it is discerned that the values in the fifth cross-validation are lower. This is attributed to the presence of more instances of leaf curling in the test set, which presents needle-like leaves, thereby increasing recognition difficulty. Apart from the challenges posed by the original shape and color of leaves, the varied backgrounds from different shooting angles also contribute to the difficulty in extracting plant leaf segments. Figure 11 provides insights into the highest and lowest segmentation scores. Specifically, it illustrates two sets of experimental data alongside their corresponding original, ground truth, and predicted images. Despite the challenging condition of leaf curling, with an IoU score of 0.6871, the segmentation results remain satisfactory.

5.2. Results of leaf withering analysis

In the analysis of leaf yellowing, the 395 sample images were categorized into three levels based on the proportion of yellow pixels among the detected leaves: less than 20% yellowing was classified as “non-withering”, 20% to 40% as “mild withering”, and over 40% as “severe withering”. Table 6 reveals sporadic errors across all categories, pri- marily due to variations in manual classification standards

and the system's reliance on RGB thresholding to detect yellow pixels, which can be influenced by environmental lighting conditions. Additionally, occlusion of leaves by flowers due to varying shooting angles contributes to errors in yellowing proportion estimation. Nonetheless, the precision, recall, and F1-score presented in Table 7 demonstrate satisfactory performance, with an overall accuracy of 96.20%, indicating the feasibility of the model in assessing plant yellowing/ withering condition.

Table 5: Performance of leaf segmentation, including mean accuracy (*mAcc*), mean intersection of unions (*mIoU*), weighted intersection of unions (*WIoU*), mean boundary F1 score (*mBF*) when employing 5-fold cross validation strategy

	mAcc	*mIoU*	*WIoU*	*mBF*
Fold 1	0.9786	0.9434	0.9901	0.9602
Fold 2	0.9577	0.9329	0.9955	0.9837
Fold 3	0.9387	0.9264	0.9889	0.9564
Fold 4	0.9612	0.9075	0.9496	0.9029
Fold 5	0.9142	0.8815	0.9492	0.8835
Average	0.9501	0.9183	0.9747	0.9373

Furthermore, Figure 10 provides the ROC curves and their corresponding AUC values for this multi-class classification task. Classes 0, 1, and 2 refer to non-, mild, and severe withering conditions. Notably, the AUC value for the non-withering class (class 0) is 0.98, while it is 0.97 for the mild withering class (class 1) and 0.98 for the severe withering class (class 2). The ROC curves for all classes remain close to the top-left corner, indicating high true positive rates and low false positive rates. Class 0, with an AUC of 0.98, shows the highest separability, suggesting that the classifier can distinguish class 0 from the other classes with exceptional accuracy. Classes 1 and 2, with AUCs of 0.97 and 0.98 respectively, also demonstrate strong performance. However, the performance for class 1 is slightly lower than that of classes 0 and 2, which could be due to the middle values being more prone to misclassification into class 0 or class 2. The substantial distance of all curves from the diagonal line, which represents random guessing (AUC = 0.5), underscores the classifier's effectiveness. These results affirm that the classifier is proficient in differentiating between the classes, making it reliable for the given multi-class classification task. Further refinement could potentially enhance the already high performance, particularly for class 1.

Table 6: Confusion matrix of plant withering analysis, where "wit." refers to withering

		Predicted		
		Non-wit.	Mild wit.	Severe wit.
Desired	Non-wit.	169	5	0
	Mild wit.	2	154	4
	Severe wit.	0	4	57

Table 7: Performance results of plant withering analysis, where "wit." refers to withering

	Accuracy	Precision	Recall	F1-score
Non-wit.	0.9713	0.9883	0.9713	0.9797
Mild wit.	0.9625	0.9448	0.9625	0.9536
Severe wit.	0.9344	0.9344	0.9344	0.9344
Average	0.9561	0.9558	0.9561	0.9559
Overall	0.9620	0.9620	0.9620	0.9620

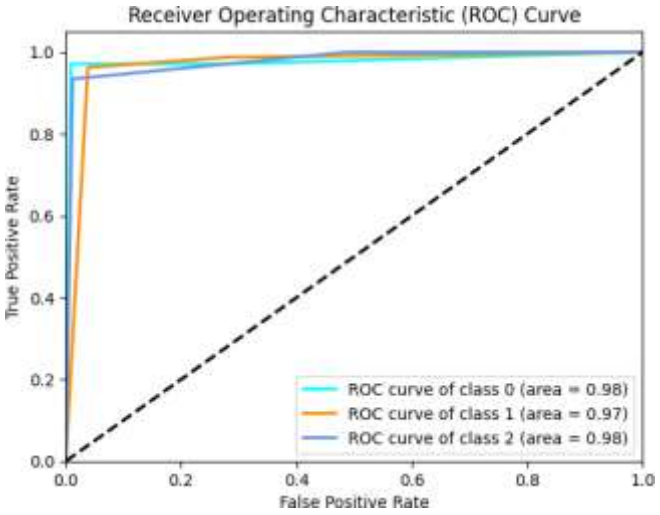

Figure 10: The ROC curves and the corresponding AUC values for multi-class classification in determining the leaf withering conditions, where classes 0, 1, and 2 refer to non-, mild, and severe withering conditions.

5.3. *Results of leaf drooping classification*

In this section, a dataset of 292 images was extracted from the database for the classification of leaf curling. It comprises 136 images depicting instances of leaf curling and 156 images without leaf curling. The dataset was divided into training and testing sets in an 8:2 ratio for cross-validation. This process was repeated five times to ensure comprehensive testing of all data. The classification results, presented in confusion matrices in Table 8, were also tabulated in Table 9 to showcase the validation out- comes across the five iterations. The average performance metrics were calculated, yielding an accuracy of 97.59%, a precision of 97.86%, a recall rate of 97.04%, and an overall F1-score of 97.24%.

Note that, although Table 9 reveals four misclassified images depicting instances of leaf curling, as illustrated in Figure 12, these misclassifications are concentrated within the same plant image. This is possibly attributed to the camera angle, which obstructs most of the leaves with flowers, making it challenging for the model to discern. Nonetheless, most of the accuracies yielded 98% and above (except the fourth fold), indicating the excellent accuracy of the model in identifying leaf curling across multiple validations.

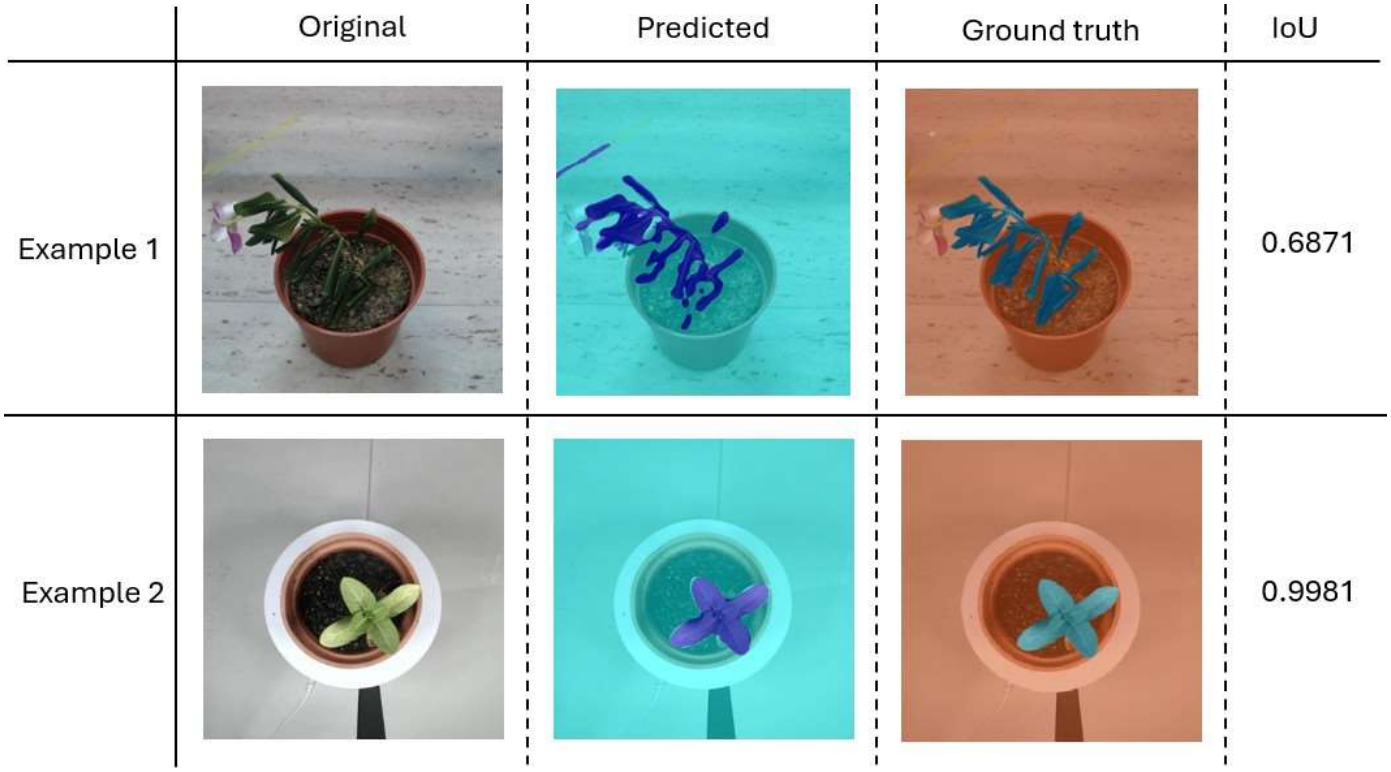

Figure 11: The example of the leaf segmentation.Top row is the case with the lowest *IoU* (i.e., 0.6871), where as bottom row is the case with the highest *IoU* (i.e., 0.9981)

Table 8: Confusion matrix of leaf curling and leaf drooping classification

		Predicted	
		Drooping	Non-drooping
Desired	Drooping	131	4
	Non-drooping	3	152

Table 9: Performance of leaf curling and drooping classification

	Accuracy	Precision	Recall	F1-score
Fold 1	0.9828	0.9643	1	0.9810
Fold 2	0.9310	1	0.8519	0.9191
Fold 3	1	1	1	1
Fold 4	0.9828	0.9643	1	0.9810
Fold 5	0.9828	0.9643	1	0.9810
Average	0.9759	0.9786	0.9704	0.9724

Figure 12: Examples of four misclassified images depicting instances of leaf curling due to the capturing angle, which obstructs most of the leaves with flowers

5.4. Results of flower counting

Similar to the previous task, a five-fold cross-validation strategy is applied for the flower counting task. Table 10 provides a detailed breakdown of the performance metrics for flower object detection for each fold. The average resultant F1-score is 85.37%. A closer inspection of each fold reveals that the fifth fold has the lowest precision

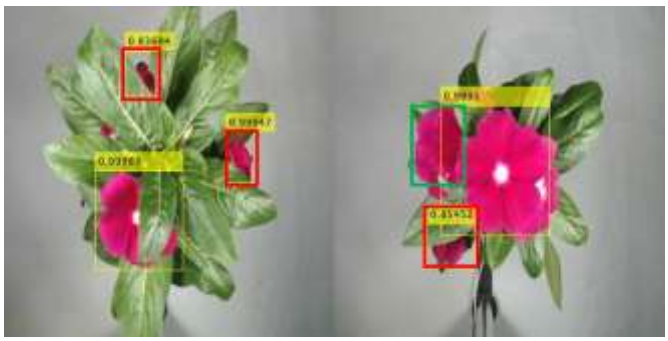

Figure 13: Examples of misdetection of the flower, resulting in incorrect flower counting, where the bounding boxes in yellow, red, and green indicate the detected flower, false positive, and false negative results, respectively

(66.18%) and consequently the lowest F1-score (78.70%). This discrepancy may be attributed to significantly different flower colors and shooting angles in the training and testing datasets. For instance, Figure 13 shows the two images with the poorest flower recognition results. Examples of flower misdetection, resulting in incorrect flower counting, are indicated by bounding boxes in yellow (detected flower), red (false positive), and green (false negative). Additionally, factors such as varying lighting conditions, background complexities, and the diversity of flowers across different geographical regions can influence the model's performance. Despite these challenges, it is reasonable to conclude that the overall performance is satisfactory and that the proposed algorithm is more robust and reliable for real-world applications.

Table 10: Performance of flower counting task in terms of precision, recall, and F1-score

	Precision	Recall	F1-score
Fold 1	0.9102	0.7988	0.8509
Fold 2	0.8316	1	0.9081
Fold 3	0.7150	0.9733	0.8244
Fold 4	0.8150	1	0.8981
Fold 5	0.6618	0.9708	0.7870
Average	0.7867	0.9486	0.8537

5.5. Summary and limitations

The proposed system, which integrates plant health diagnostics and an irrigation system, offers several significant strengths. Firstly, its use of advanced CNN models for tasks such as leaf segmentation, leaf yellowing/withering identification, leaf curling classification, and flower counting ensures high accuracy and robustness across various plant conditions and environmental factors. This consistent performance, as depicted in the histogram in Figure 14, highlights the robustness and reliability of the proposed system. The experimental results demonstrate that the system consistently achieves performancelevels exceeding 75% across all tasks, with certain metrics, such as *mAcc* and *mIoU*, reaching over 90%. This high level of accuracy is crucial for precise plant health monitoring and decision-making, enabling timely interventions and optimized resource utilization. Additionally, the system's ability to process and analyze data from multiple angles further enhances its reliability, making it well-suited for real-world agricultural applications.

Moreover, the system's integration of plant health diagnostics with automated irrigation offers a comprehensive solution for precision agriculture. By accurately detecting signs of withering, drooping, and flowering, the system can automate the irrigation process, ensuring plants receive the appropriate amount of water and nutrients based ontheir specific needs. This not only improves plant health and productivity but also promotes sustainable farming practices by minimizing water and fertilizer waste. Theuse of a newly collected dataset tailored to various growth conditions and angles also highlights the system's adaptability and potential for widespread application in diverse agricultural settings.

Despite its strengths, the proposed system does have certain limitations. One notable challenge is the system's sensitivity to variations in environmental conditions, such as lighting and background complexities, which can affect the accuracy of the diagnostic algorithms. For instance, the flower counting task exhibited discrepancies due to differences in flower colors and shooting angles between the training and testing datasets. Additionally, the presence of occlusions, such as flowers obstructing leaves, posed challenges for accurate leaf condition assessment. These factors highlight the need for further refinement and calibration to ensure consistent performance across different scenarios.

Furthermore, while the system demonstrates high accuracy overall, there are instances where specific tasks, such as leaf curling classification, resulted in misclassifications due to obstructed views and challenging angles. The fifth fold in the cross-validation process, which included more instances of curled leaves, showed lower performance metrics, indicating areas where the model could be improved. Addressing these limitations will require ongoing research and development to enhance the system's robustness and adaptability, ensuring it can reliably handle the diverse and dynamic conditions encountered in real-world agricultural environments.

6. Conclusion

In summary, this work introduces a novel approach and framework for plant monitoring and care, integrating advanced deep learning image recognition with a quad-axis mechanical slider system. The proposed system, capable of bidirectional communication with MATLAB, effectively assesses plant growth conditions through various analytical methods. Training was conducted using a dataset comprising 926 images that capture diverse growth conditions of plants. Utilizing semantic segmentation, object

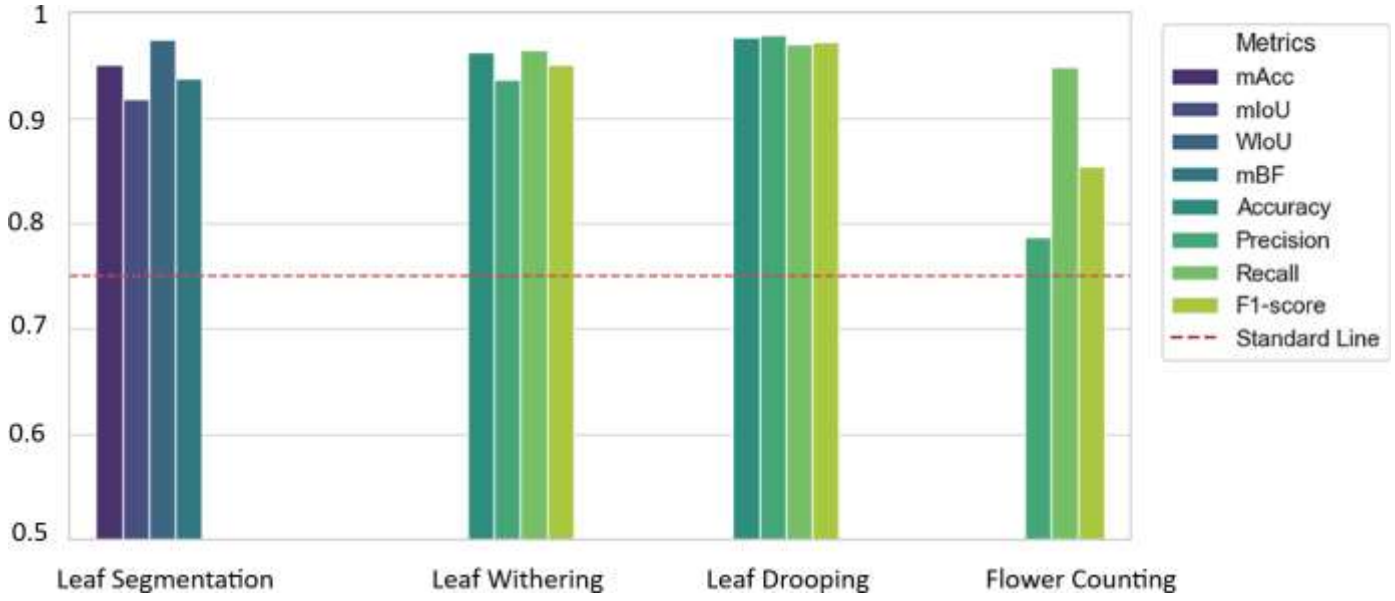

Figure 14: Overall performance on the four tasks: leaf segmentation, leaf withering analysis, leaf drooping classification, and flower counting. A standard line is drawn at 0.75 as the baseline indicator for all tasks.

classification, and object detection methods, the system is proficient in performing tasks such as leaf yellowing analysis, determination of leaf curling, and flower counting. Extensive analysis and investigation confirm the system's reliability and robustness, achieving an average accuracy of approximately 95% for leaf segmentation, 96.20% overall accuracy in leaf withering analysis, 97.59% accuracy in leaf drooping classification, and an F1-score of 85.37% for the flower counting task, thereby validating the effectiveness of the collected database.

The proposed system's high accuracy and comprehensive functionality demonstrate its potential as a valuable tool in precision agriculture. By providing real-time mon- itoring and analysis of plant health, the system can facil- itate more informed decision-making, optimize the use of resources like water and fertilizers, and ultimately enhance crop yield and quality. Its adaptability to various growth conditions and angles underscores its practicality for real- world agricultural applications. The integration of plant health diagnostics with automated irrigation exemplifiesa holistic approach to plant care, promoting sustainable farming practices by minimizing resource wastage.

This study should incentivize future research to address the limitations identified, such as the system's sensitivity to environmental conditions and the challenges posed by occluded views and varied backgrounds. Enhancing the robustness and adaptability of the diagnostic algorithmsto handle a wider range of conditions will be a priority. Additionally, expanding the dataset to include more di-verse plant species and growth stages will help improve the model's generalizability. The incorporation of advanced sensor technologies and IoT (Internet of Things) capabilities could further enhance the system's precision and functionality, enabling more comprehensive plant health monitoring and automated care solutions. Continued research and development will aim to refine the system for broader application in diverse agricultural settings, contributing to the advancement of precision agriculture.

Acknowledgments

This work was funded by Ministry of Science and Technology (MOST) (Grant Number: MOST 111-2221-E-035-059-MY3).

References

[1] M. O. Agriculture, Inquiry, https://agrstat.moa.gov.tw/sdweb/public/inquiry/InquireAdvance.aspx, retrieved: 2024-05-29 (2024).
[2] A. Pal, Ai pal, https://www.aipaltech.com/, retrieved: 2024-05-29 (2024).
[3] M. Toğaçar, B. Ergen, Z. Cömert, Classification of flower species by using features extracted from the intersection of feature selection methods in convolutional neural network models, Measurement 158 (2020) 107703.
[4] M. Cıbuk, U. Budak, Y. Guo, M. C. Ince, A. Sengur, Effi- cient deep features selections and classification for flower species recognition, Measurement 137 (2019) 7–13.
[5] C. Zheng, T. Liu, A. Abd-Elrahman, V. M. Whitaker, B. Wilkinson, Object-detection from multi-view remote sensing images: A case study of fruit and flower detection and count- ing on a central florida strawberry farm, International Journal of Applied Earth Observation and Geoinformation 123 (2023) 103457.
[6] D. Li, J. Li, S. Xiang, A. Pan, Psegnet: Simultaneous seman- tic and instance segmentation for point clouds of plants, Plant Phenomics (2022).
[7] S. P. Mohanty, D. P. Hughes, M. Salathé, Using deep learn- ing for image-based plant disease detection, Frontiers in plant science 7 (2016) 1419.

[8] D. E. Goldberg, J. H. Holland, Genetic algorithms and machine learning, Machine Learning 3 (1988) 95–99. doi:10.1023/A:1022602019183.
[9] F. Glover, Future paths for integer programming and links to artificial intelligence, Computers & operations research 13 (5) (1986) 533–549.
[10] C. Cortes, V. Vapnik, Support-vector networks, Machine learning 20 (1995) 273–297.
[11] K. Simonyan, A. Zisserman, Very deep convolutional net- works for large-scale image recognition, arXiv preprint arXiv:1409.1556 (2014).
[12] A. Krizhevsky, I. Sutskever, G. E. Hinton, Imagenet classification with deep convolutional neural networks, Advances in neural information processing systems 25 (2012) 1097–1105.
[13] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, L. Fei-Fei, Ima- genet: A large-scale hierarchical image database, in: 2009 IEEE conference on computer vision and pattern recognition, Ieee, 2009, pp. 248–255.
[14] D. Hughes, M. Salathé, et al., An open access repository of images on plant health to enable the development of mobile disease diagnostics, arXiv preprint arXiv:1511.08060 (2015).
[15] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, A. Rabinovich, Going deeper with convolutions, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2015, pp. 1–9.
[16] A. Conn, U. V. Pedmale, J. Chory, S. Navlakha, High-resolution laser scanning reveals plant architectures that reflect universal network design principles, Cell systems 5 (1) (2017) 53–62.
[17] A. Conn, U. V. Pedmale, J. Chory, C. F. Stevens, S. Navlakha, A statistical description of plant shoot architecture, Current biology 27 (14) (2017) 2078–2088.
[18] R. Girshick, Fast r-cnn, in: Proceedings of the IEEE interna- tional conference on computer vision, 2015, pp. 1440–1448.
[19] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2016, pp. 770–778.
[20] F. Schroff, D. Kalenichenko, J. Philbin, Facenet: A unified embedding for face recognition and clustering, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2015, pp. 815–823.
[21] M. Ester, H.-P. Kriegel, J. Sander, X. Xu, et al., A density-based algorithm for discovering clusters in large spatial databases with noise, in: kdd, Vol. 96, 1996, pp. 226–231.
[22] L.-C. Chen, G. Papandreou, F. Schroff, H. Adam, Rethink- ing atrous convolution for semantic image segmentation, arXiv preprint arXiv:1706.05587 (2017).
[23] S. Ren, K. He, R. Girshick, J. Sun, Faster r-cnn: Towards real- time object detection with region proposal networks, Advancesin neural information processing systems 28 (2015).