

Hydra: Accurate Multi-Modal Leaf Wetness Sensing with mm-Wave and Camera Fusion

Yimeng Liu¹, Maolin Gan¹, Huaili Zeng¹, Li Liu², Younsuk Dong¹, Zhichao Cao^{1*}

¹Michigan State University, USA ²Tsinghua University, China

Abstract

Leaf Wetness Duration (LWD), the time that water remains on leaf surfaces, is crucial in the development of plant diseases. Existing LWD detection lacks standardized measurement techniques, and variations across different plant characteristics limit its effectiveness. Prior research proposes diverse approaches, but they fail to measure real natural leaves directly and lack resilience in various environmental conditions. This reduces the precision and robustness, revealing a notable practical application and effectiveness gap in real-world agricultural settings. This paper presents Hydra, an innovative approach that integrates millimeter-wave (mm-Wave) radar with camera technology to detect leaf wetness by determining if there is water on the leaf. We can measure the time to determine the LWD based on this detection. Firstly, we design a Convolutional Neural Network (CNN) to selectively fuse multiple mm-Wave depth images with an RGB image to generate multiple feature images. Then, we develop a transformer-based encoder to capture the inherent connection among the multiple feature images to generate a feature map, which is further fed to a classifier for detection. Moreover, we augment the dataset during training to generalize our model. Implemented using a frequency-modulated continuous-wave (FMCW) radar within the 76 to 81 GHz band, Hydra's performance is meticulously evaluated on plants, demonstrating the potential to classify leaf wetness with up to 96% accuracy across varying scenarios. Deploying Hydra in the farm, including rainy, dawn, or poorly light nights, it still achieves an accuracy rate of around 90%.

CCS Concepts

• **Computer systems organization** → **Sensors and actuators**; • **Applied computing** → **Agriculture**.

*Zhichao Cao is the corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

ACM MobiCom '24, September 30–October 4, 2024, Washington D.C., USA

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0489-5/24/09.

<https://doi.org/10.1145/3636534.3690662>

Keywords

Agricultural IoT, Precision Farming, Multi-Modality Sensing

ACM Reference Format:

Yimeng Liu¹, Maolin Gan¹, Huaili Zeng¹, Li Liu², Younsuk Dong¹, Zhichao Cao¹. 2024. Hydra: Accurate Multi-Modal Leaf Wetness Sensing with mm-Wave and Camera Fusion. In *International Conference On Mobile Computing And Networking (ACM MobiCom '24)*, November 18–November 22, 2024, Washington D.C., DC, USA. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3636534.3690662>

1 Introduction

Agriculture plays a crucial role in the global economy, contributing 4% to the worldwide GDP and over 25% in some developing countries [48]. However, plant disease's increasing frequency and severity threaten productivity, food security, and biodiversity, particularly in vulnerable regions [43, 51]. One critical factor influencing plant disease spread is leaf wetness, which is water on leaf surfaces, occurs due to dew, precipitation, fog, or irrigation and signifies [33, 50]. The duration of the leaf wetness is an essential factor for the growth of pathogens such as *Venturia inaequalis* [20, 53]. Thus, accurate LWD detection is vital for effective disease control [41]. Numerous agricultural studies have demonstrated that accurate LWD sensing can help protect yields across various crops, such as strawberries [30], grapes [14], and lettuce [56].

Therefore, researchers have introduced different modalities to enhance the performance of the LWD detection system [9, 10, 13, 34, 49]. Nevertheless, existing systems face inherent limitations:

Wetness detection on real leaf: The Leaf Wetness Sensor (LWS) employs synthetic leaves [13, 34, 49], which vary from actual leaves in size, shape, and material, leading to a detection error of up to 30 minutes, as detailed in Section 6.1.

Environment Resilience: The real environment, dynamic and complex nature due to varying light, wind, and plant diversity, can significantly affect existing detection methods. For instance, RGB imaging techniques [9] are highly sensitive to lighting conditions, and mm-Wave-based approach [10] struggle with leaf vibrations from the wind. Additionally, the diverse characteristics of plants, such as shape and size, pose substantial challenges for deploying sensing systems.

System efficiency: Frequent detection is crucial for accurate disease management, especially on large farms where wetness changes across many plants must be tracked multiple



Figure 1: Leaf Wetness detection with Hydra.

times per hour. This demand for frequent sensing highlights the need for system efficiency to ensure that more plants can be monitored regularly. However, compared to traditional cameras and sensors, the mm-Wave-based approach takes longer to create an RF image by scanning plants [10]. Therefore, improving system efficiency is essential to maintain high accuracy in disease management.

As illustrated in Table 1, previous works either fail to operate on real leaves [13, 34, 49] effectively or lack efficiency [10]. They are also sensitive to environmental changes, resulting in inconsistent performance [9, 10].

This paper introduces Hydra, as shown in Figure 1, a multi-modal system combining RGB camera and mm-Wave radar to detect whether the leaf status is wet or dry, achieving high accuracy and efficiency in leaf wetness detection. The basic observation is that RGB camera and mm-Wave radar are two orthogonal channels. RGB images are not sensitive to leaf vibration, and mm-Wave radar is resilient to various light conditions. Cameras and mm-Wave radar generate imaging, which is the input of a Deep Neural Network (DNN) based fusion model. It includes some layers to enhance its system resilience without model retraining. To our knowledge, Hydra is the pioneering system to combine RGB and mm-Wave modalities for leaf sensing. It excels over previous systems in its environmental robustness, high efficiency, and capability for deployment on real leaves.

However, several challenges must be addressed in developing Hydra:

i) **How to fusion the mm-Wave and Camera image:** We apply an FMCW chirp, sensitive to varying distances from the mm-Wave radar, to image the plant. This technique produces multiple 2D images representing the plant’s cross-sections at different depths, indicating the distance from the cross-sections to the radar. Meanwhile, RGB cameras generate a singular image that captures surface details influenced

Table 1: Comparison Methods for Detecting LWD

Methods	Real Leaf	System Efficiency	Environment Resilience
LWS [13, 34, 49]	No	Yes	No
Infrared/RGB cameras [9]	Yes	Yes	No
mm-Wave radar [10]	Yes	No	No
Hydra	Yes	Yes	Yes

by lighting conditions. Developing a method that effectively combines a mm-Wave image at a certain depth with the high-resolution surface detail from RGB images is crucial for creating a comprehensive view of plant wetness. However, due to the different ways of representing leaf spatial information, it is challenging to fuse mm-Wave and camera to utilize their orthogonal information. To address this challenge, we design a single-depth modality fusion and feature extraction method that takes multiple mm-Wave images at different depths and an RGB image as inputs. Specifically, for each depth, we selectively mask the mm-Wave image at that depth with the RGB image to generate a fusion depth image. Then, we design a CNN model to extract multiple features from these fusion depth images at different depths.

ii) **How to combine the fusion features at a 3D level:** Integrating the fusion features at different depths necessitates the development of a robust 3D classification model capable of accurately interpreting complex, multi-dimensional information. However, since the 3D structure of various crops is usually complex and challenging to model, finding a way to connect the features at different depths as a representative feature map is not trivial. To solve this problem, we develop a transformer-based feature encoder, which treats the fusion features at different depths as a serial sequence and automatically finds the inherent 3D domain connection among them to generate an informative feature map.

iii) **How to increase the system efficiency of the mm-Wave imaging and keep the high accuracy:** mm-Wave technology produces high-resolution images through its short wavelength and SAR technique, proving essential for detailed observations. However, the scanning period significantly degrades the system efficiency. Enhancing system efficiency while preserving high-resolution imaging in mm-Wave technology is complex due to the compromise between synthetic aperture size and detail accuracy. Finding ways to boost efficiency without sacrificing the quality of leaf wetness detection remains a significant challenge. To tackle this challenge, we observe that our two-modality fusion enhances the feature domain, allowing us to shrink mm-Wave Radar’s Field-of-View (FoV). We design an adaptive FoV adjustment method to balance the system’s efficiency and accuracy.

iv) **How to adapt the system in different scenarios:** In real-world conditions, plants are subjected to constantly changing environments that can significantly impact the

performance of leaf wetness detection. Therefore, it is imperative to enhance Hydra’s robustness, ensuring it can accurately detect and analyze leaf wetness under a wide range of environmental conditions. This adaptation involves creating algorithms resilient to light, temperature, and movement changes, ensuring reliable detection across diverse agricultural settings. Addressing these challenges is vital for advancing precision agriculture and contributing to more effective plant disease management and crop production strategies. To address this challenge, we augment the datasets used for model training. Specifically, we include different lighting conditions for camera images. For mm-Wave radar, we manually generate diverse leaf vibration patterns with a fan for different leaf types, sizes, gestures, etc.

We implement Hydra with commercial off-the-shelf hardware components. By applying our model to various plants in different environmental scenarios, we have successfully trained a deep-learning model notable for its accuracy and robustness. Extensive testing demonstrates that Hydra attains an outstanding accuracy rate of 96.63% under indoor conditions with simulated scenarios and maintains a performance level of about 90% in the farm, including rainy conditions, dawn, and poorly nighttime settings. Moreover, Hydra exhibits exceptional performance by reducing the error margin in LWD detection to approximately 2 minutes, markedly outperforming previous systems.

The contributions of this paper can be summarized as follows:

- To our knowledge, Hydra is the first contactless, multi-modality sensing system designed explicitly for LWD detection. It can directly scan plant surfaces for accurate, robust, and efficient wetness identification.
- We have incorporated novel multi-modality techniques, leveraging the synergistic capabilities of mm-Wave imaging and camera image that allow for an advanced data fusion, significantly mitigating the challenges of aligning disparate dimensional data.
- We improved the SAR system’s efficiency by 25%, enabling precise leaf wetness detection. Combined with RGB imaging, this enhancement allows for accurate assessments within a limited field of view and using a smaller synthetic aperture. This optimization notably enhances the system’s ability to measure moisture levels accurately.
- The system demonstrates robust performance in complex outdoor scenarios, maintaining high efficiency even when specific modalities may not function optimally. When our model is applied to the farm, it achieves around 90% accuracy, showcasing its adaptability and effectiveness in varied environmental settings.

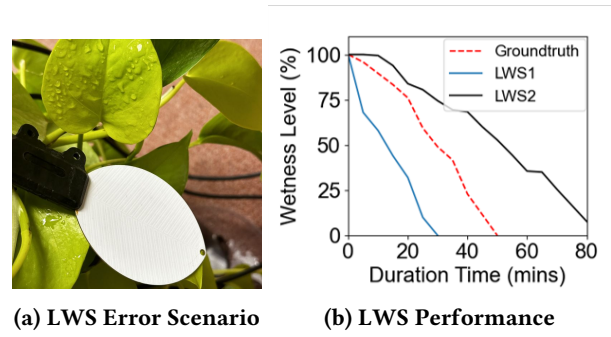


Figure 2: LWS Analysis. Leaf wetness error scenario demonstrating the inaccuracy of LWS [13].

2 Preliminary and Motivation

In this section, we introduce the definition of Leaf Wetness and Leaf Wetness Duration and review prior efforts in LWD. We begin with the existing LWS and the RGB camera solution. We then discuss approaches that utilize mm-Wave radar and outline the motivation behind our research.

2.1 Definition of Leaf Wetness

Leaf Wetness Duration is the period during which a leaf remains wet. Accurately detecting LWD is crucial for understanding and managing plant diseases [41]. A vital part of determining LWD is accurately identifying Leaf Wetness, which refers to the presence of water on a leaf’s surface [33, 50]. This condition can be classified into two states: dry or wet. In our study, we obtain the ground-truth leaf wetness using a moisture meter specified in Section 5.3. We calibrate the meter in a dry, ventilated indoor environment to establish a baseline threshold. If the moisture reading exceeds this threshold, the leaf is considered wet. Otherwise, the leaf is dry.

2.2 Leaf Wetness Sensor

The Leaf Wetness Sensor [13, 34, 49] mimics the leaf’s shape and positioning to simulate the real leaf status. Within this domain, LWSs are categorized into resistive, using hydrophilic to detect wetness through resistance changes [49], and capacitive types, leveraging dielectric material coatings to detect wetness via capacitance shifts, based on the differing dielectric constants of water and air [13]. The Bio-Mimetic LWS [34] uses the capacitive LWS to simulate the surface and texture of leaves to detect LWD more accurately. These LWSs can represent the change of the free water on the leaf. However, a major issue with these sensors is that they do not directly measure the wetness of actual leaves, so they may not precisely show the result.

We conduct an experiment using two different Capacitive LWS [13] placed at various locations on the plant to

Table 2: Camera Performance in Different Lighting

Lighting Condition	Normal Light	Dawn	Evening
Camera Accuracy	90.42%	72.73%	68.97%

detect changes in wetness. We record the wetness level by normalizing the LWS readings to a 0-100 scale, in which 0 represents complete dryness and 100 represents saturated wetness. The ground truth detection is based on the wetness meter described in Section 5.3. We record and normalize the wetness meter readings to a 0-100 scale. Our observation in Figure 2a indicates that the sensors typically weigh more than actual leaves, causing them to slide more easily and leading to quicker water runoff than real leaves. Additionally, in specific scenarios, the sensor’s edge texture may trap water, preventing evaporation and resulting in prolonged wetness duration compared to natural leaf surfaces. We conduct an empirical study to illustrate the difference in wetness level between the ground truth and LWSs in the same environment. The analysis in Figure 2b reveals that the sensors varying placements on the plant lead to inconsistent results, significantly deviating from the established ground truth. This limitation underscores the need for higher fidelity in LWD detection and contributes to more precise and reliable agricultural monitoring systems.

2.3 RGB Camera Recognition

Exploring LWD through RGB cameras offers clear distinctions between dry and wet leaves, with obvious features with drops on the leaf that can be effectively detected with DNN under ideal lighting conditions [35]. However, the accuracy significantly drops in agriculture, where lighting varies—during dawn, night, or cloudy days. The performance of only a camera for detection across various lighting conditions decreased. We empirically study the influence of light conditions on RGB image-based detection. We are using ResNet-18 [16] to classify if the leaf is dry or wet with the RGB image. The results are summarized in Table 2. The result in normal light is around 90%, but when the lighting condition degrades, it drops to 70%. Also, RGB cameras rely on reflected light, limiting their ability to detect conditions behind or beneath overlapping leaves. They cannot penetrate the plant’s canopy to reveal hidden wetness.

2.4 mm-Wave-based Sensing

mm-Wave, operating at the higher frequencies of the electromagnetic spectrum, leverages the advantage of large bandwidth at shorter distances [55]. This characteristic limits penetration and propagation depth in underwater environments, as high-frequency electromagnetic (EM) signals have limited penetration depth [37]. Conversely, this limitation benefits

Table 3: SAR Image Performance in different scenario.


Wind Condition	No Wind	Moderated Wind
SAR Image Accuracy	90.85%	70.43%

surface texture applications [26]. The precision and sensitivity afforded by mm-Wave technology, mainly through the dramatic drop of the reflection, enable the detection of minute wetness variations on leaf surfaces.

Synthetic Aperture Radar (SAR) is widely used in radar applications for its ability to produce high-resolution images by simulating a large aperture through the relative motion between the radar and the target. A recent study demonstrated the use of SAR in enhancing imaging resolution and environmental perception for autonomous systems in real-time conditions [61]. The increased aperture size allows SAR to capture fine-grained images, revealing intricate details that improve leaf wetness detection and enable precise monitoring and accurate identification of leaf status. Using the FMCW chirp, mm-Wave radar is sensitive to the distance between the radar and the target. SAR can image targets at different depths, providing cross-sectional information at varying distances from the radar.

From the methodology described by [58], we use a testbed incorporating the mm-Wave to apply the SAR technique imaging a potted plant, capturing the reflected signals from the leaves. This approach detects leaf wetness by directly targeting the plant, offering a sensitive solution that avoids mere sampling. In our adaptation of [10], we position a stationary plant approximately 200 mm from the radar, monitoring the transition from wet to dry conditions. Observations are recorded at 20-minute intervals, with the plant reaching complete dryness after 80 minutes. Figure 3a illustrates an example of a SAR imaging result in 230mm depth, a cross-section of a plant 230 mm away from the radar. The image features an area highlighted in yellow, encircled by a red rectangle, designating the leaf surface information detected by the radar. The irregular shape of plants makes it difficult to identify the specific part of the plant depicted in an image and to detect features related to wetness. The analysis shown in Figure 3b suggests that within a depth range of 200-340mm, the leaf surface area reflection signals from the plant increase as it dries. This observation highlights the potential of using reflection intensity as a feature of plant dryness, offering insights into plants’ wetness content.

Current methods have notable limitations, especially in terms of their robustness. The accuracy predominantly depends on analyzing data from a constrained depth range, which inadequately addresses leaf wetness and plant health comprehensively. Apart from that, the environment poses a significant challenge for mm-Wave imaging. We use the ResNet-18 model to classify the SAR image between dry and



(a) SAR Imaging (b) Wetness Stage vs Reflection Coefficient

Figure 3: Example of SAR imaging results and the signal reflection change as the leaf dries.

wet. The results in Table 3 show that wind significantly impacts SAR imaging accuracy. In conditions with no wind, the accuracy can reach 90%, but with moderate wind, the precision drastically drops to 70%. This finding highlights the challenges of SAR imaging in windy conditions.

2.5 Motivation

We have observed that no existing single-modality leaf wetness sensing techniques can achieve high accuracy under various environmental conditions. This motivates us to consider a multi-modality solution to enable high-accuracy leaf wetness sensing with strong environmental resilience while keeping system efficiency as high as possible.

3 System Overview

We present Hydra, an advanced multi-modal leaf wetness detection system. Hydra integrates mm-Wave and RGB imaging to facilitate water on the leaf wetness detection across various plant species. Then, we use the detection result to calculate LWD to help with disease control. Its precise and durable system architecture, shown in Figure 4, consists of the main components: SAR Imaging System, Single-Depth Feature Extraction, Multi-Depth Feature Detection, Model Training and Data Enhancement.

SAR Imaging System. In this module, we design the SAR imaging system that transmits mm-Wave signals and senses the target plants from various depths. Our approach leverages the foundational design principles outlined in [10] to optimize efficiency by incorporating an RGB camera.

Single-Depth Feature Extraction. In this module, Hydra system introduces a groundbreaking fusion method that integrates data from mm-Wave radar images captured in a certain depth with RGB camera images containing surface information. By merging these diverse modalities with different ways of representing leaf spatial information, we generate a comprehensive input set that enhances the accuracy of subsequent detection processes.

Multi-Depth Feature Detection. In this module, our approach employs an innovative fusion method that consolidates data across multiple depths for the final classification outcome. This strategy develops a thorough insight into the plant’s characteristics, culminating in a holistic understanding that informs the overall result.

Model Training and Data Enhancement. This module covers our strategies for assessing and improving the model’s performance under various environmental conditions, aimed at increasing its robustness and versatility for complex agricultural scenarios.

4 Hydra Design

This section explores the methodologies behind Hydra, beginning with an overview of SAR imaging techniques and proceeding to the extraction of single-depth features. It then delves into the detection of leaf wetness using multi-depth feature analysis. Finally, we discuss the model training protocols and data enhancement strategies employed.

4.1 Variable Scan Distance SAR Imaging

In this section, we develop our imaging system based on SAR technology. This system emits signals at regular intervals to capture images of relatively stationary targets. One key strength of SAR technology is its ability to increase the aperture size effectively. With the combination of FMCW and SAR, imaging technology has the advantages of being lightweight and low-cost and is widely used in near-field imaging scenarios [31, 45, 57, 58].

For the FMCW signal, the instantaneous frequency can be expressed as a linear function as

$$m(t) = \cos[2\pi(f_0 t + 0.5Kt^2)] \quad (1)$$

where f_0 represents the carrier frequency at the initial time $t = 0$, and $K = B/T$ denotes the frequency modulation slope, derived from the sweep bandwidth B and the chirp duration T . Upon receiving the backscattered signal, the radar system employs a dechirping process that isolates the beat frequency by mixing the received signal with its in-phase ($s_I(t)$) and quadrature ($s_Q(t)$) components. This yields a complex beat signal expressed as

$$s(t) = s_I(t) - js_Q(t) = \sigma e^{-j2\pi(f_0\tau + K\tau t - 0.5K\tau^2)}, \quad (2)$$

with τ signifying the round-trip delay of the echo, and σ encompassing both the target’s reflectivity and amplitude decay.

For the spatial representation of the received signal, particularly in the wave number domain, the signal can be simplified as

$$s(x', y_T, y_R, k) = \iiint p(x, y, z) e^{-jkR_T} e^{-jkR_R} dx dy dz, \quad (3)$$

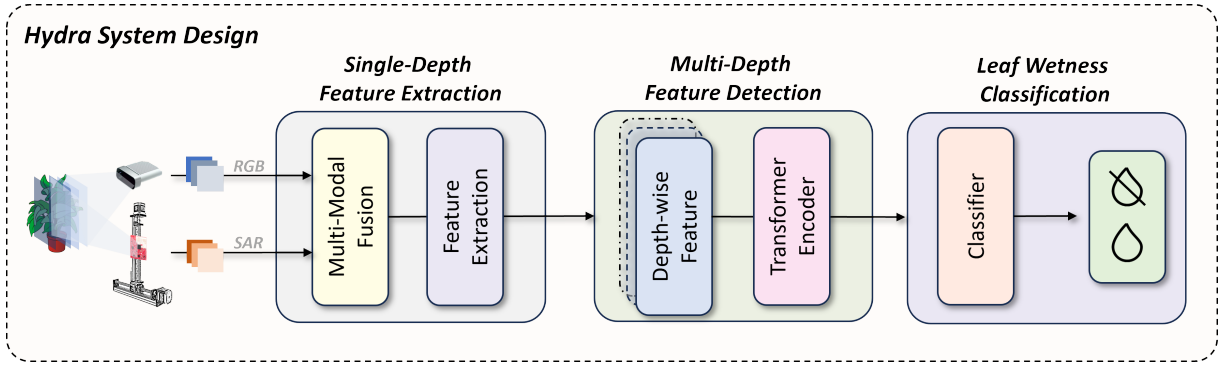


Figure 4: Hydra Overview. Hydra contains three main procedures: **Single-Depth Feature Extraction**, **Multi-Depth Feature Detection**, and classifier. We fuse the multi-modality and extract wetness features in the **Single-Depth Feature Extraction** phase. The **Multi-Depth Feature Detection** stage leverages these features at a 3D level, culminating in utilizing a leaf wetness classification algorithm to derive the final wetness assessment.

where R_T and R_R denote the distances from the transmitter and receiver to the scatter point, respectively, aligning within a Cartesian coordinate system where x , y , and z axes define the scanning, vertical, and depth dimensions.

To accurately model the distances R_T and R_R about the transmitter and receiver's positions, we consider

$$\begin{aligned} R_T &= \sqrt{(x + (x' + \delta_T/2))^2 + (y - y_T)^2 + (z - Z_0)^2}, \\ R_R &= \sqrt{(x - (x' + \delta_T/2))^2 + (y - y_T)^2 + (z - Z_0)^2}. \end{aligned} \quad (4)$$

This geometric consideration is vital for the subsequent image reconstruction phase, where the range migration algorithm is adapted for the SAR context. The conversion from multi-static to a monostatic equivalent involves phase compensation, simplifying the data processing for precise depth-based imaging as

$$\tilde{s}(x', y') = \iint p(x, y) e^{-j2kR} dx dy, \quad (5)$$

highlighting the distance R to the scatter point and facilitating the extraction of 2D images from the radar data.

Finally, incorporating Weyl's representation theorem [54] enables the approximation of spherical waves through plane waves, leading to the backscatter data being expressed as

$$\tilde{S}(k_x, k_y) = P(k_x, k_y) e^{jk_z Z_0}, \quad (6)$$

where the inverse Fourier transform applied to $\tilde{S}(k_x, k_y)$ yields the reconstructed 2D image of the target area to get the 2D image:

$$p(x, y) = IFT^{(k_x, k_y, k_z)} [e^{-jk_z Z_0} \tilde{S}(k_x, k_y)], \quad (7)$$

Adjusting the x range balances a broader field-of-view and higher resolution versus a shorter view for enhanced efficiency. This enables flexibility in capturing detailed imagery or prioritizing rapid scanning based on specific requirements.

4.2 Single-Depth Feature Extraction

Integrating SAR imaging with camera images poses a significant challenge due to their different ways of representing leaf spatial information. Enhanced by FMCW chirps, SAR imaging excels at detecting distance variations, allowing for the extraction of detailed features at multiple depths. This capability includes identifying features not visible to the RGB camera. Conversely, camera imaging relies on light reflection to capture information, project, and record data based on light reflected off the scanning plane. This method can capture details from multiple depth layers within a single image. This difference complicates the direct integration of the two image types, as they provide information in distinct dimensions.

To tackle this challenge, we introduce a novel approach incorporating a machine learning layer designed to fuse these two modalities, as illustrated in Figure 5. We employ the concept of a mask for dimension alignment, using a certain depth SAR image to mask the camera image. We've adapted this conventional computer vision masking technique to spotlight regions of interest. We normalize the SAR image to a $[0, 1]$ scale, where higher values indicate pixels of greater relevance. This normalized mask is then applied to the camera image to extract the information from the depth specified by the SAR image. Moreover, the fusion process preserves the SAR image's ability to reveal surface textures through reflection rates. We also introduce a dynamic parameter within the layer to adjust the balance between the SAR and camera images based on feature correspondence.

Following the data fusion, the feature detection process becomes significantly more complex. In this context, CNNs emerge as an invaluable tool for our analysis. CNNs excel in navigating the complexities of high-dimensional data, adeptly isolating and enhancing features critical for accurate

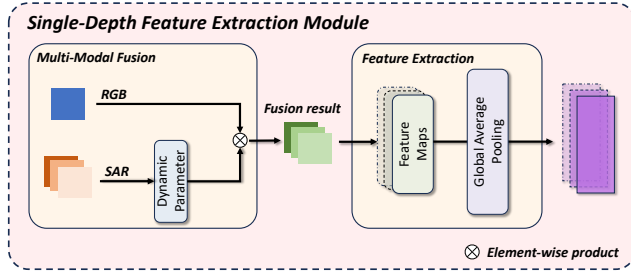


Figure 5: Single-Depth Feature Extraction architecture.

classification. The approach of mimicking the hierarchical pattern recognition found in the human visual system enables the layered structure to progressively refine inputs, starting with essential edge detection and advancing to more complex feature identification [44].

In the concluding phase of our Single-Depth Feature Extraction module, we introduce a Global Average Pooling (GAP) layer [27]. This layer calculates the spatial mean of feature maps from the terminal convolutional layer, streamlining data complexity while preserving critical feature information without adding new trainable parameters. The GAP layer is a structural regularizer. It reduces overfitting by simplifying the architecture and improving the model’s interpretability. It can enable the application of Class Activation Mapping (CAM) to produce heatmaps to show the model’s focus areas. This feature is handy when analyzing SAR images and the outputs from the multi-modality fusion module, where direct feature visualization is challenging. From Figure 6, we can get the example of RGB image, SAR image in various depths, and CAM. Following the fusion of multiple modalities, this example shows that the CAM heatmap reveals the model’s focus on leaf areas vital for wetness evaluation, corroborated by both the camera and SAR imagery.

To enhance the model’s receptivity to detailed SAR features, we selectively exclude camera data from the fusion process and focus solely on SAR imaging. We de-emphasized RGB imagery in training by randomly removing 20% of the images from the dataset. This approach encourages the model to develop a deeper understanding of SAR features and achieve a well-rounded grasp of multi-modal data. By employing this method, we significantly boosted the model’s ability to identify subtle wetness cues on leaves.

4.3 Multi-Depth Leaf Wetness Detection

A notable challenge in plant analysis is the complexity of plant structures. Varying leaf sizes, orientations, and densities make accurate interpretation difficult. While 3D models provide a detailed view and can improve the accuracy of the analysis. We have developed a model by integrating and classifying data across multiple depths, enabling practical 3D




Figure 6: Example of Class Activation Mapping Result.

analysis shown in Figure 7. We first introduce depth-aware positional encoding [11], which is meticulously designed to incorporate depth positional information into our model and help spatial positioning across depth dimensions. The mathematical underpinning of this layer is expressed through the positional encoding formula:

$$PE(depth, i) = \begin{cases} \sin\left(\frac{depth}{10000^{2i/d_{model}}}\right) & \text{for even } i \\ \cos\left(\frac{depth}{10000^{2(i-1)/d_{model}}}\right) & \text{for odd } i \end{cases} \quad (8)$$

Here, $depth$ signifies the position within the depth sequence, i represents the dimension within the feature space, and d_{model} indicates the feature space dimensionality. Then, we implement a multi-head attention mechanism reminiscent of those in cutting-edge Large Language Models (LLMs), which endows our transformer encoder with exceptional contextual awareness [2–4]. This feature supports comprehensive analysis throughout the plant by merging data from various depths, offering a detailed and precise assessment of leaf. The multi-head attention mechanism concurrently focuses on multiple depth-related aspects to pinpoint subtle moisture shifts. Integrating depth-aware positional encoding and multi-head attention mechanism improves our model’s precision in leaf wetness detection by addressing the complexities of spatial analysis and positional data synthesis.

4.4 Model Training and Data Enhancement

Training the model for leaf wetness detection involves a meticulously designed process to accurately classify leaves’ wetness and dryness states. The foundation of our model’s training regimen is based on utilizing binary cross-entropy as the loss function. This decision is based on the effectiveness of binary cross-entropy for binary classification tasks, which makes it especially apt for differentiating between the binary states of leaf surface conditions. The loss function can be shown as

$$L(y, \hat{y}) = -\frac{1}{N} \sum_{i=1}^N [y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)], \quad (9)$$




Figure 7: The Multi-Depth leaf wetness Detection module's architecture.

where L is the loss, N is the number of observations, y_i represents the actual label for the i^{th} observation, and \hat{y}_i denotes the predicted probability of the leaf being wet. Our model's training begins with the pretraining phase of the Single Depth Fusion Module, an essential element for accurately detecting leaf wetness across varying depths. We will apply Section 4.3 with the extracted feature to combine the feature analysis and produce the final classification results. Moreover, we introduce simulations of breeze-induced motion, a condition under which SAR imaging may not faithfully capture the scene. We train our model to recognize and adjust to such anomalies. This holistic data augmentation strategy, addressing multiple real-world scenarios, significantly enhances the model's resilience and relevance to environmental conditions.

5 Implementation

5.1 System Implementation

We developed a prototype of Hydra¹ As shown in Figure 8, the SAR Imaging system features a two-axis mechanical scanner designed to enable high-speed acquisition. The scanner's design parameters are finely tuned for optimal plant analysis. It has a horizontal traversal capacity of 150 mm and a vertical range of 100 mm, aligning with the expected synthesis aperture size and the dimensions of our target plant. The radar on the scanning mechanism is a Texas Instruments (TI) IWR1642 [47], operating within the 77 to 81 GHz frequency range. This radar unit is essential for capturing the raw mmWave signals reflected from the targets under study. Complementing the radar, a DCA 1000EVM [21] is employed for initial signal collection and processing, setting the stage for detailed imaging analysis. In addition to the SAR Imaging system, our prototype includes a camera imaging component, leveraging an Azure Kinect camera [32], which is placed at the center of the scanning area. This camera captures reflections in the visible light spectrum, thus providing a multimodal imaging perspective. Integrating camera imaging with mmWave radar-based SAR imaging offers a holistic

¹The source design are available at <https://github.com/liuyime2/MobiCom24-Hydra>.

view of the target, significantly enhancing the analysis's depth and accuracy. In evaluating the precision performance of our models, we've adopted the rigorous approach of applying 5-fold cross-validation, iterated ten times across different datasets. This method is crucial for reducing the potential variability and bias that might arise from the dataset's diversity, ensuring that our performance metrics accurately reflect the model's capability under various conditions.

5.2 System Preparation

Data Calibration. In the system design outlined in Section 4.2, the successful fusion of SAR and RGB camera images hinges on aligning their FoV precisely. This alignment is crucial for integrating these images effectively in subsequent training processes. Given the complexity of the match-up, calibration between SAR imaging and camera imaging systems becomes a significant task. To address this problem, we employ a novel calibration strategy using a stable, depth-consistent object: the scissor. This object is chosen because it can remain at a fixed depth, making it an ideal calibration target. The calibration process is illustrated with the SAR imaging results of the scissor displayed in Figure 9. We precisely match the RGB camera image, highlighted by the red block in Figure 9b.

Feature Extractor. In Section 4.2, we detail a sophisticated feature extractor developed to process complex inputs comprising camera and SAR images, laying the groundwork for the subsequent Multi-Depth Leaf Wetness Detection module introduced in Section 4.3. We have explored the utility of pre-trained CNN models, utilizing a benchmark dataset ImageNet [7], renowned for its extensive range of image classes tailored for classification tasks. Our investigation encompasses leading models like ResNet [16], VGG [42], DenseNet [19], and InceptionV3 [46], each acclaimed for their architectural innovations and exceptional performance in image classification challenges. Our comparative analysis, presented in Table 4, reveals that ResNet surpasses its counterparts in accuracy for our specific application. The exceptional performance of ResNet is primarily attributed to its revolutionary architecture. It integrates residual learning to overcome the vanishing gradient problem [17], a common challenge in DNN. This design is particularly effective for processing the complex, high-dimensional data associated with SAR imaging, thereby enhancing leaf wetness's feature extraction and classification process amidst the diverse and complex plant foliage shapes.

5.3 Data Collection

Ground truth. Without a dedicated sensor to accurately track changes in LWd, we have adopted a novel approach to establishing reliable ground truth. We employ a commercial




Figure 8: Hydra comprises three components: a mmWave radar, a two-axis scan testbed, and an RGB camera.




Figure 9: Calibration of the RGB camera and SAR sensor with the scissor. The red rectangle is the corresponding field of view for the MIMO sensor.

moisture meter from General Tools [12], traditionally used for measuring wood surface wetness levels. This device operates on the principle of electrical resistance measurement between two pins, as depicted in the accompanying Figure 10. Although not initially designed for leaf surface wetness detection, this instrument offers a viable alternative for our purposes. By gently tapping the moisture meter’s pins onto the leaf surface, we can observe variations in electrical resistance that correlate with the leaf’s moisture content; higher water presence increases the readings. This method allows us to measure the wetness level on the leaf surface directly. Based on these resistance readings, we can effectively establish a threshold to differentiate between various leaf dryness and wetness stages.

Baseline. In our experimental setup, we initially employed the commercial LWS, specifically the PHYTOS 31 [13], to monitor leaf wetness. We use the PHYTOS 31 default threshold for the readings to determine whether the leaf is wet or dry. We faced a significant challenge with the sensor’s metal components, which enhanced backscatter and potentially affected the SAR imaging system’s data quality. Additionally, these components obstructed light, creating a barrier for the RGB camera. To mitigate these issues, we selected an analogous plant of similar size and placed it within the same environment as the monitored specimen. To ensure comprehensive coverage and accurate assessment of leaf wetness across the plant, we positioned four different sensors, distributing them evenly to capture the entirety of the plant’s

Table 4: Comparison Methods for Detecting Wetness

Pre-trained Model	SAR Image Acc	Camera Acc
DenseNet	87.58% ± 6.10%	83.62% ± 1.97%
VGG-16	85.50% ± 4.31%	82.42% ± 1.83%
InceptionV3	89.96% ± 5.69%	83.16% ± 2.13%
ResNet-18	92.61% ± 3.43%	85.85% ± 1.31%

foliage. Given the inherent variation in baseline dryness resistance across different sensors, we focused on analyzing the average incremental change in resistance to determine leaf wetness. This method bypasses the variability of individual sensor readings, leveraging the collective data from multiple sensors to provide a more reliable and consistent measure of LWD. We also incorporate baseline comparisons for camera [35] and mmLeaf [10] technologies to ensure a comprehensive approach to wetness detection. We selected ResNet-18 as the optimal model for camera-based wetness detection based on comparisons in Table 4. For the mmLeaf system, we followed the procedures in [10] to align our scanning and modeling with best practices for wetness detection. **Data Collection.** Our indoor experiment involved a six-month data collection period from five diverse plant types, focusing on leaf size and orientation variations. During this period, the plants exhibited different growth patterns, spacing, and distributions, increasing the dataset’s diversity. We used the same plants in outdoor experiments to assess performance in dynamic environments. Additionally, we conducted in-situ experiments on two fields totaling 3.67 acres, planted with soybeans and corn, to evaluate Hydra in practical settings. The experiments were performed under various scenarios, including different times of the day—early morning, afternoon, and evening—and diverse weather conditions, such as sunny, windy, and post-rain environments. The experiment is set up with the plant 200-500 mm away from the mmWave radar and focusing on two extreme values: completely dry and saturated wet. This approach maximizes the range of wetness levels, enabling us to distinguish wetness features effectively. We compiled a dataset of approximately 674 pairs, including 536 pairs from indoor environments, 46




Figure 10: Using Moisture Meter for the ground truth Measurement.

pairs from outdoor environments, and 92 pairs from real fields with dynamic environments. Additionally, we developed a dataset for Hydra to evaluate the performance of the LWD. Hydra monitor plants transitioning from fully saturated to completely dry indoors. We sample the entire drying process 6 to 12 times. Twenty groups are collected from varying plants and environmental conditions depending on the drying time. This diverse collection helped highlight the differences in wetness features on leaf surfaces. Leveraging the refined resolution capabilities of mmWave technology, we captured images at 1 mm depth intervals up to a distance of 300 mm, ensuring comprehensive coverage of plant information. Additionally, to construct a precise timeline of the drying process, we monitored the changes in leaf wetness every 10 minutes, establishing a ground truth for drying duration. Our evaluation also encompasses outdoor scenarios, covering sunny and windy conditions with temperatures between 10-17 Celsius degrees, spanning from morning to night, to consider the potential for dew formation. Given Hydra’s limitation of not operating in the rain, we simulated rainy conditions for the plant alone to assess the system’s adaptability and performance under varied environmental influences.

6 EVALUATION

This section details the performance evaluation of Hydra through comprehensive over-the-air experiments, showcasing its capabilities in real-world scenarios.

6.1 Overall Performance

In our evaluation of Hydra, we focused on its precision in discerning various levels of leaf surface wetness. Utilizing an extensive dataset, our system demonstrated an exceptional ability to identify, achieving a median precision rate of $96.21\% \pm 2.57\%$, as illustrated in Figure 11a. This performance was compared against traditional methods, camera imaging, the mmLeaf system, and LWS systems, where Hydra outperformed with accuracies of $87.8\% \pm 2.32\%$, $83.84\% \pm 3.21\%$ and $75.25\% \pm 4.43\%$, respectively.

In the LWD analysis, a comparative test across 20 different data groups as detailed in Section 5.3 pitted Hydra against conventional methods: LWS, Camera, and mmLeaf systems. The result shown in Figure 11b demonstrates Hydra’s accuracy, detecting LWD accurately within a 2-minute error margin in 18 out of 20 groups. In comparison, the Camera system’s performance varied, with eight groups having 5-10-minute errors and two groups showing 15-20-minute errors. The mmLeaf system saw four groups with 5-10-minute errors and two with 15-20-minute errors. LWS results were the most varied, with errors spanning from 5 minutes to 30 minutes across different groups, highlighting Hydra’s precision in detecting the LWD.

These results highlight the advancements Hydra brings to the field by addressing the limitations of other modalities. Camera, as discussed in Section 2.3, lacks depth information, making it difficult to detect wetness when leaves overlap. The mmLeaf system’s effectiveness is compromised by a lack of environmental resilience, as noted in Section 2.4. LWS accuracy heavily depends on precise sensor placement. The differences in weight and surface texture can affect wetness dynamics, leading to lower accuracy and higher variability, as stated in Section 2.2. Hydra’s comprehensive approach combines depth information and surface texture analysis, overcoming these limitations.

6.2 Single Depth Fusion Performance

In Hydra, we introduced a single depth fusion designed to merge data from camera image and SAR imaging, targeting accurately identifying leaf wetness features. To evaluate the effectiveness of our single-depth fusion approach, we conducted a comparative study against established data fusion techniques, specifically early fusion and late fusion. Early fusion combines input data at the model’s onset, while late fusion integrates features before the classification stage. Additionally, our analysis compares our unique data enhancement strategies, as detailed in Section 4.4. In our evaluation, we compared our single-depth fusion module, which integrates RGB camera and SAR imaging data, against traditional early and late fusion methods. Our analysis, detailed in Figure 12a, highlighted that conventional methods struggle with the dimensional differences between modalities, limiting detection accuracy to $83.42\% \pm 1.84\%$ and $82.83\% \pm 1.47\%$ for early and late fusion, respectively. In contrast, our fusion module significantly outperformed these, achieving an accuracy of $90.37\% \pm 2.24\%$. Data augmentation techniques, including removing and adjusting RGB camera images, enhanced our module’s precision to $93.02\% \pm 1.79\%$. The results demonstrate our module’s ability to utilize multimodal data effectively. Data enhancement leads to more stable results, effectively allowing the model to uncover deeper features




Figure 11: Hydra performance of detection LWD.

within both SAR and RGB images, which can reduce the impact of overfitting.

6.3 Scan Motion Performance

As mentioned in Section 1, efficiency is crucial for large farm applications. Increased efficiency allows farms to monitor more plants with Hydra, leading to more accurate disease control. This section evaluates scanning duration and FoV to optimize our system’s efficacy without compromising accuracy. In our evaluation, we trained the model using a single type of plant fixed at 200 mm from the mm-Wave sensor under optimal lighting, focusing exclusively on assessing the effects of scan distance modifications and comparing them with mmLeaf. Through a series of experiments, we systematically adjusted the scanning field from a comprehensive 200 mm, ensuring entire plant coverage, down to a focused 100 mm, centering on the plant’s core. As shown in Figure 12b, for Hydra, our findings establish a direct correlation between the width of the scanning field and detection accuracy. With a maximal field of 200mm, the system achieved its highest accuracy when reducing the FoV, which led to a steady decline in accuracy. For the scan distance longer or equal to 150 mm, with the advantage of the multi-modality, the accuracy is above 90%, which is $95.43\% \pm 1.47\%$, $94.68\% \pm 1.97\%$ and $93.38\% \pm 2.56\%$ for 200mm, 175mm and 150mm respectively. For the scan distance lower than 150mm, the accuracy reduces dramatically with a higher variance of $89.28\% \pm 4.05\%$ and $84.1\% \pm 4.53\%$. The degradation in performance can be attributed to a decrease in image resolution. Additionally, a shorter scan distance causes Hydra to miss critical parts necessary for accurate detection. For the mmLeaf, the accuracy drops rapidly from 200 mm to 100 mm; the accuracy is $89.66\% \pm 1.46\%$, $85.71\% \pm 2.79\%$, $84.83\% \pm 2.63\%$, $82.18\% \pm 3.27\%$, $81.13\% \pm 3.83\%$. This trend underscores that reducing the scan distance leads to a smaller synthetic aperture and narrower FoV, which diminishes the detection of wetness features on leaves. However, leveraging multimodal features can compensate for this loss, ensuring the SAR, compared

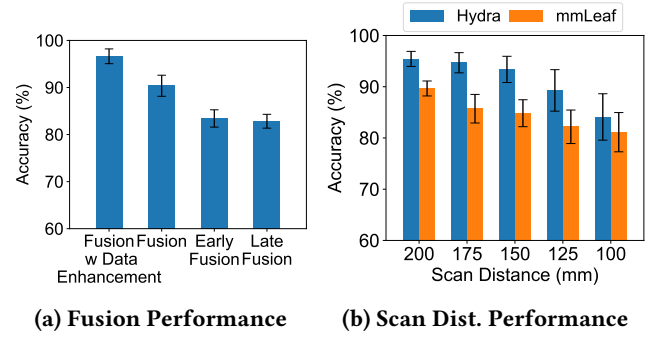


Figure 12: Performance of Multi-Modal fusion and SAR system’s performance across scan distances

to mmLeaf, has a shorter efficiency of 25% and retains high accuracy by enhancing the overall picture of leaf wetness through integrating diverse data sources.

6.4 Application in Different Environment

We evaluated sensors’ performance in the indoor environment and simulated real farm environments as shown in Figure 14, focusing on how environmental factors impact the accuracy of wetness detection technologies relative to the baseline. Our study involves three sensors: Hydra, a standard camera, and mmLeaf. We examined their precision in detecting leaf wetness while considering the analyzed environmental variables.

We evaluate the wind environment under a controlled indoor environment using a fan to simulate the impact of the wind. The result is illustrated in Figure 13a. Without wind, accuracy rates were high across all devices: 95.76% for Hydra, 90.85% for the camera, and 89.2% for mmLeaf. The introduction of moderate wind saw a decline, particularly affecting Hydra and mmLeaf due to their reliance on SAR imaging, where movement disrupts phase alignment, leading to errors. Despite this, Hydra maintained a leading precision rate of 90.62%. The camera’s quicker imaging speed has little influence on achieving 88.2%. But mmLeaf dropped to 70.43%, showcasing the wind has significantly influenced the SAR imaging process. However, Hydra maintains its robustness in wind scenarios thanks to its multi-modal features and a training process designed for resilience.

We evaluate Hydra performance for the different types of leaves in the indoor environment, which focuses on the leaf orientation and size. For larger leaves positioned leaf-side down, Hydra, camera, and mmLeaf showcased exceptional accuracy, registering at 95.76%, 92.61% and 88.22% as shown in Figure 13b. Hydra precision experienced a minor decline to 93.53% when the system was tasked with evaluating smaller leaves oriented leaf-side up and a dramatic degradation for

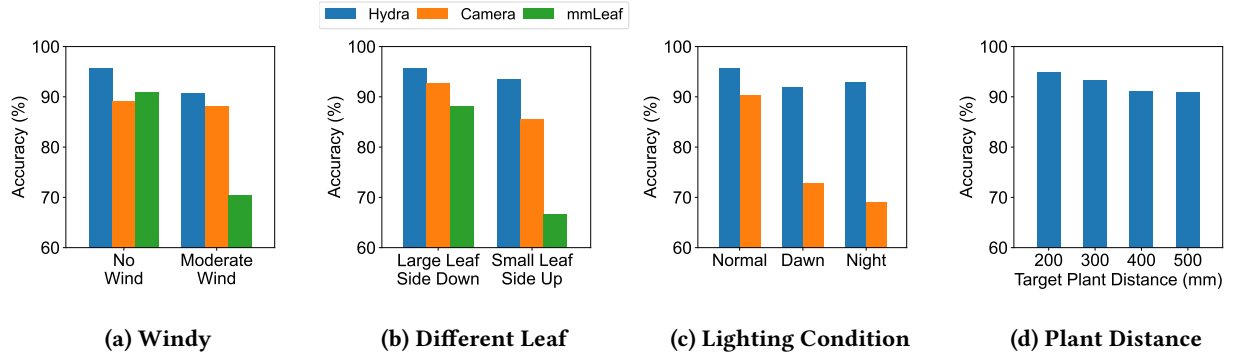


Figure 13: Performance for Hydra application in a different real-world scenario.



Figure 14: Hydra application in different scenarios.

camera and mmLeaf of the precision 85.5%, 66.57%. The precision shows that Hydra is robustness to the plant with different morphology. The precision of Hydra demonstrates its robust ability to adapt to plants with varying morphologies, ensuring accurate wetness detection across different plant types.

We evaluate the impact of different lighting conditions from normal to dawn to night in the outdoor simulated real farm environment. The analysis, summarized in Figure 13c, demonstrates that lighting significantly influences camera-based detection, with accuracy dropping in low-light conditions. In regular light, the camera can reach an accuracy of 90.42%, and at Dawn and night, it will drop to 72.73% and 68.97%. However, Hydra exhibited minimal impact from varying light conditions, underscoring its comprehensive multimodal analysis capability and resilience in diverse environmental settings. The accuracy would drop a little but still keep the accurate precision, which the normal light, dawn time, and night time are 95.63%, 91.87%, 92.92%, respectively.

Further, evaluate Hydra system assessed its performance across different distances from a target plant in an indoor setting as shown in Figure 13d. The highest accuracy, 94.86%, was achieved when the plant was 200mm away. As the distance increased to 300mm, 400mm, and 500mm, accuracy slightly decreased to 93.23%, 90.08%, and 91.94%, respectively,

indicating a relationship between increased distance and diminished mm-Wave resolution. The results suggest that while Hydra with the multi-modal maintains robust performance across varying distances, optimizing placement relative to the target can enhance detection precision.

6.5 Application in Real Field

We conducted real-field experiments in soybean and corn fields, including early morning, dawn, evening, sunny, and rainy weather. Our goal was to examine Hydra precision in dynamic, practical environments.

We first simulated practical conditions outdoors and then tested in soybeans and corn fields. The results, shown in Figure 15a, indicate that the simulated outdoor environment and real farm conditions achieved similar accuracies of 89.13% and 89.6%, respectively. Although these are slightly lower than the indoor setup accuracy of 94.64%, they still demonstrate superior performance.

Soybeans and corn are two distinct types of crops that differ significantly in leaf size and arrangement. Soybeans have smaller, denser leaves, while corn features more extensive, sparsely arranged ones. The results for Hydra are shown in Figure 15b. We achieved a high accuracy rate of 88.89% for corn and 90.2% for soybeans. Evaluating diverse leaf structures with high accuracy demonstrates Hydra robustness in detecting wetness across various crop types.

The environment in a real farm is inherently complex due to its ever-changing lighting and weather conditions. Lighting varies significantly throughout the day, from soft morning light to intense afternoon sunlight to dim evening light. These changes affect the visibility and reflectivity of leaves, impacting the accuracy of Hydra. Figure 15c shows an accuracy of 90.91% in the morning, 88% at dawn, and 88.89% in the evening. The relatively lower accuracy during dawn suggests that direct sunlight might mislead the camera. In the morning, Hydra benefits from optimal multimodality conditions, while in the evening, it relies mainly on mmWave




Figure 15: Performance for Hydra application in a different real-world scenario.

information. Additionally, weather conditions such as sunny and rainy periods further complicate the environment. We observed an accuracy of 90.2% during sunny conditions and 88.46% during rainy conditions, as shown in Figure 15d. Results from evaluations at different times of day and various weather conditions demonstrate Hydra’s robustness in real farm settings.

7 RELATED WORK

Leaf Water Content. Innovations in wireless technology for monitoring plant moisture include methods like RFID, backscatter, and PCR technologies, primarily designed to measure Leaf Water Content (LWC) [6, 8, 18]. Additionally, technologies such as LoRa [23, 38–40], ultrasound [24], and infrared sensing [59] offer non-destructive and non-invasive means to accurately assess LWC.

mm-Wave Multimodality. Recent advancements in mm-Wave and multimodal technologies have enhanced health diagnostics and wireless communication. Studies [1, 29] demonstrate improved gait recognition and health monitoring accuracy through mm-Wave imaging and micro-Doppler fusion. Additionally, multimodal features enhance model training, object detection, depth estimation, and mm-Wave beam selection in wireless networks [5, 25, 52].

3D Classification. The Hydra model uses 3D data analysis to tackle complex classification challenges, enhancing precision in identifying intricate patterns beyond 2D capabilities. In healthcare, Hydra improves diagnostic accuracy in tumor identification and anatomical mapping, supporting personalized treatment plans. It has been applied to detect moldy corn [60] and analyze sleep postures for better healthcare outcomes [28].

8 PRACTICAL DEPLOYMENT DISCUSSION

Deploying Hydra in agriculture enhances monitoring and disease control, leveraging advanced technology to optimize practices and increase crop yields in large farms and greenhouses.

Drones are increasingly used for precision agriculture because they cover large areas efficiently. For large farms, Hydra can be mounted on drones to autonomously navigate predefined routes, providing real-time data collection across extensive fields. This approach is supported by studies on drone applications in crop monitoring and pesticide spraying [15, 22].

In greenhouses or dense fields where drones may struggle with leaf wetness detection due to space constraints and thick foliage, Hydra can utilize rail systems for continuous monitoring [36]. These rail-mounted systems efficiently navigate greenhouses, offering targeted analysis by sampling selected plants. This method allows precise data collection and real-time environmental adjustments, optimizing plant health and growth in confined or densely planted areas.

9 CONCLUSION

This paper introduces the development, deployment, and comprehensive evaluation of Hydra, an advanced system for precise leaf wetness detection. Hydra tackles the critical challenge of accurately assessing leaf wetness, a significant factor in agricultural productivity and disease management. We introduced Hydra to incorporate and analyze data from multiple sources, including innovative SAR imaging and an RGB camera. We conducted extensive experiments to evaluate Hydra in practical environments. Our experiments included various plants and environmental factors. The results show that Hydra consistently surpasses existing methodologies in accuracy and reliability. Hydra demonstrated a remarkable 96% accuracy in determining leaf wetness across diverse settings. When deployed on a real farm, Hydra consistently maintained around 90% accuracy.

Acknowledgement

We sincerely thank the anonymous reviewers and our shepherd for their valuable feedback. This work was partially supported by NSF CAREER Award 2338976.

References

- [1] Abdulwahab Alazeb, Mouazma Batool, Naif Al Mudawi, Mohammed S. Alshehri, Sultan Almakdi, Nouf Abdullah Almajally, and Asaad Algarni. 2024. Effective Gait Abnormality Detection in Parkinson's Patients for Multi-Sensors Surveillance System. *IEEE Access* 12 (2024), 48686–48698. <https://doi.org/10.1109/ACCESS.2024.3372593>
- [2] Bocheng Chen, Nikolay Ivanov, Guangjing Wang, and Qiben Yan. 2023. DynamicFL: Balancing Communication Dynamics and Client Manipulation for Federated Learning. In *Proceedings of IEEE SECON*.
- [3] Bocheng Chen, Advait Paliwal, and Qiben Yan. 2023. Jailbreaker in Jail: Moving Target Defense for Large Language Models. In *Proceedings of ACM MTD*.
- [4] Bocheng Chen, Guangjing Wang, Hanqing Guo, Yuanda Wang, and Qiben Yan. 2023. Understanding Multi-Turn Toxic Behaviors in Open-Domain Chatbots. In *Proceedings of ACM RAID*.
- [5] Z Chen, R Tian, D Xiong, C Yuan, T Li, and Y Shi. 2024. Multi-Dimensional Information Fusion You Only Look Once Network for Suspicious Object Detection in Millimeter Wave Images. *MDPI Electronics* 13, 4 (2024), 773. <https://www.mdpi.com/2079-9292/13/4/773>
- [6] Spyridon Nektarios Daskalakis, George Goussetis, Stylianos D Assimonis, Manos M Tentzeris, and Apostolos Georgiadis. 2018. A uW backscatter-morse-leaf sensor for low-power agricultural wireless sensor networks. *IEEE Sensors Journal* 18, 19 (2018), 7889–7898.
- [7] J. Deng, W. Dong, R. Socher, L. Li, Kai Li, and Li Fei-Fei. 2009. ImageNet: A large-scale hierarchical image database. In *Proceedings IEEE CVPR Workshops*.
- [8] Shuvashis Dey, Emran Md Amin, and Nemai Chandra Karmakar. 2020. Paper based chipless RFID leaf wetness detector for plant health monitoring. *IEEE Access* 8 (2020), 191986–191996.
- [9] S Duvdevani. 1947. An optical method of dew estimation. *Quarterly Journal of the Royal Meteorological Society* 73, 317-318 (1947), 282–296.
- [10] Maolin Gan, Yimeng Liu, Li Liu, Chenshu Wu, Younsuk Dong, Huacheng Zeng, and Zhichao Cao. 2023. Poster: mmLeaf: Versatile Leaf Wetness Detection via mmWave Sensing. In *Proceedings of ACM MobiSys*.
- [11] Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N. Dauphin. 2017. Convolutional sequence to sequence learning. In *Proceedings of ACM ICML*.
- [12] General Tool. 2024. Moisture Meter. <https://generaltools.com/digital-tools/moisture-humidity-digital-tools/moisture-meters>. Accessed: 2024-03-15.
- [13] METER Group. 2021. *PHYTOS 31 Manual Web*. METER Group. Retrieved Nov 21, 2022 from http://library.metergroup.com/Manuals/20434_PHYTOS31_Manual_Web.pdf
- [14] WD Gubler, MR Rademacher, SJ Vasquez, and CS Thomas. 1999. Control of powdery mildew using the UC Davis powdery mildew risk index. *APSnet Feature. Published online. The American Phytopathological Society, St. Paul, MN* (1999).
- [15] Abdul Hafeez, Mohammed Aslam Husain, S.P. Singh, Anurag Chauhan, Mohd. Tauseef Khan, Navneet Kumar, Abhishek Chauhan, and S.K. Soni. 2023. Implementation of drone technology for farm monitoring & pesticide spraying: A review. *Information Processing in Agriculture* 10, 2 (2023), 192–203. <https://doi.org/10.1016/j.inpa.2022.02.002>
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *IEEE CVPR*.
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *Proceedings IEEE CVPR*.
- [18] NA Hoog, TE van den Berg, and HS Bindra. 2022. A 60 GHz pulsed coherent radar for online monitoring of the withering condition of leaves. *Sensors and Actuators A: Physical* 343 (2022), 113693.
- [19] Gao Huang, Zhuang Liu, and Kilian Q. Weinberger. 2016. Densely Connected Convolutional Networks. *Proceedings of IEEE CVPR* (2016).
- [20] L Huber and TJ Gillespie. 1992. Modeling leaf wetness in relation to plant disease epidemiology. *Annual review of phytopathology* 30, 1 (1992), 553–577.
- [21] Texas Instruments. 2024. DCA1000EVM. <https://www.ti.com/tool/DCA1000EVM>. Accessed: 2024-03-15.
- [22] Mohamed Lachgar, Hamid Hrimech, Ali Kartit, and Brahim Hlimi. 2023. Unmanned aerial vehicle-based applications in smart farming: A systematic review. *International Journal of Advanced Computer Science and Applications* 14, 6 (2023), 819–831. <https://doi.org/10.14569/IJACSA.2023.0140696>
- [23] Chenning Li, Yidong Ren, Shuai Tong, Shakhrol Iman Siam, Mi Zhang, Jiliang Wang, Yunhao Liu, and Zhichao Cao. 2024. ChirpTransformer: Versatile LoRa Encoding for Low-power Wide-area IoT. In *Proceedings of ACM MobiSys*.
- [24] Gen Li, Zhichao Cao, and Tianxing Li. 2023. EchoAttack: Practical Inaudible Attacks To Smart Earbuds. In *Proceedings of ACM MobiSys*.
- [25] H Li, Y Ma, Y Huang, Y Gu, W Xu, Y Liu, and Others. 2024. RIDERS: Radar-Infrared Depth Estimation for Robust Sensing. *arXiv preprint arXiv:2402.02067* (2024). <https://arxiv.org/abs/2402.02067>
- [26] Yumeng Liang, Anfu Zhou, Huanhuan Zhang, Xinzhe Wen, and Huadong Ma. 2021. FG-LiquidID: A Contact-less Fine-grained Liquid Identifier by Pushing the Limits of Millimeter-wave Sensing. *Proceedings of ACM UbiCom* (2021).
- [27] Min Lin, Qiang Chen, and Shuicheng Yan. 2014. Network In Network. *arXiv:1312.4400 [cs.NE]*
- [28] Chen Liu, Zixuan Dong, Li Huang, Wenlong Yan, Xin Wang, Dingyi Fang, and Xiaojiang Chen. 2024. TagSleep3D: RF-based 3D Sleep Posture Skeleton Recognition. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 8, 1 (2024). <https://doi.org/10.1145/3643512>
- [29] C Ma and Z Liu. 2024. mDS-PCGR: A Bi-Modal Gait Recognition Framework With the Fusion of 4D Radar Point Cloud Sequences and micro-Doppler Signatures. *IEEE Sensors Journal* (2024). <https://ieeexplore.ieee.org/abstract/document/10413972/>
- [30] SJ MacKenzie and NA Peres. 2012. Use of leaf wetness and temperature to time fungicide applications to control anthracnose fruit rot of strawberry in Florida. *Plant disease* 96, 4 (2012), 522–528.
- [31] Adriano Meta, Peter Hoogeboom, and Leo P Ligthart. 2007. Signal processing for FMCW SAR. *IEEE Transactions on Geoscience and Remote Sensing* 45, 11 (2007), 3519–3532.
- [32] Microsoft. 2024. Azure Kinect DK. <https://azure.microsoft.com/en-us/products/kinect-dk>. Accessed: 2024-03-15.
- [33] Concepció Moragrega and Isidre Llorente. 2023. Effects of Leaf Wetness Duration, Temperature, and Host Phenological Stage on Infection of Walnut by *Xanthomonas arboricola* pv. juglandis. *MDPI Plants* 12, 15 (2023). <https://doi.org/10.3390/plants12152800>
- [34] Brian H Nguyen, Gregory S Gilbert, and Marco Rolandi. 2023. A Bio-Mimetic Leaf Wetness Sensor from Replica Molding of Leaves. *Advanced Sensor Research* 2, 6 (2023), 2200033.
- [35] Arth Patel, Won Suk Lee, Natalia A Peres, and Clyde W Fraisse. 2021. Strawberry plant wetness detection using computer vision and deep learning. *Smart Agricultural Technology* 1 (2021), 100013.
- [36] Vivek Ramakant Pathmudi, Narendra Khatri, Sandeep Kumar, Antar Shaddad Hamed Abdul-Qawy, and Ajay Kumar Vyas. 2023. A systematic review of IoT technologies and their constituents for smart and sustainable agriculture applications. *Scientific African* 19 (2023), e01577.
- [37] Umair Mujtaba Qureshi, Faisal Karim Shaikh, Zuneera Aziz, Syed M. Zafi S. Shah, Adil A. Sheikh, Emad Felemban, and Saad Bin Qaisar. 2016. RF Path and Absorption Loss Estimation for Underwater Wireless Sensor Networks in Different Water Environments. *MDPI Sensors* 16,

- 6 (2016). <https://doi.org/10.3390/s16060890>
- [38] Yidong Ren, Puyu Cai, Jinyan Jiang, Jialuo Du, and Zhichao Cao. 2023. Prism: High-throughput LoRa Backscatter with Non-linear Chirps. In *Proceedings IEEE INFOCOM*.
- [39] Yidong Ren, Li Liu, Chenning Li, Zhichao Cao, and Shigang Chen. 2022. Is LoRaWAN Really Wide? Fine-grained LoRa Link-level Measurement in An Urban Environment. In *Proceedings IEEE ICNP*.
- [40] Yidong Ren, Wei Sun, Jialuo Du, Huaili Zeng, Younsuk Dong, Mi Zhang, Shigang Chen, Yunhao Liu, Tianxing Li, and Zhichao Cao. 2024. Demeter: Reliable Cross-soil LPWAN with Low-cost Signal Polarization Alignment. In *Proceedings of ACM MobiCom*.
- [41] Tracy Rowlandson, Mark Gleason, Paulo Sentelhas, Terry Gillespie, Carla Thomas, and Brian Hornbuckle. 2015. Reconsidering leaf wetness duration determination for plant disease management. *Plant Disease* 99, 3 (2015), 310–319.
- [42] Karen Simonyan and Andrew Zisserman. 2014. Very Deep Convolutional Networks for Large-Scale Image Recognition. *CoRR* abs/1409.1556 (2014). <https://api.semanticscholar.org/CorpusID:14124313>
- [43] Brajesh K Singh, Manuel Delgado-Baquerizo, Eleonora Egidi, Emilio Guirado, Jan E Leach, Hongwei Liu, and Pankaj Trivedi. 2023. Climate change impacts on plant pathogens, food security and paths forward. *Nature Reviews Microbiology* 21, 10 (2023), 640–656.
- [44] Ghislain St-Yves, Emily J. Allen, Yihan Wu, Kendrick Kay, and Thomas Naselaris. 2023. Brain-optimized deep neural network models of human visual areas learn non-hierarchical representations. *Nature Communications* 14, 1 (Jun 2023). <https://doi.org/10.1038/s41467-023-38674-4>
- [45] Andrew G Stove. 1992. Linear FMCW radar techniques. In *Proceedings IET Radar and Signal Processing*.
- [46] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna. 2015. Rethinking the Inception Architecture for Computer Vision. arXiv:1512.00567 [cs.CV]
- [47] Texas Instruments. 2024. IWR1642. <https://www.ti.com/product/IWR1642>. Accessed: 2024-03-15.
- [48] THE WORLD BANK. 2024. Agriculture and Food. <https://www.worldbank.org/en/topic/agriculture/overview>. Accessed: 2024-03-14.
- [49] THE WORLD BANK. 2024. Leaf Wetness Sensor, Vantage Pro2™ and EnviroMonitor. <https://www.davisinstruments.com/products/leaf-wetness-sensor-vantage-pro-and-vantage-pro2>. Accessed: 2024-03-15.
- [50] UC SANTA CRUZ. 2024. Leaf wetness sensor will enable better plant disease forecasting and management. <https://news.ucsc.edu/2023/03/leaf-wetness-sensors.html>. Accessed: 2024-03-15.
- [51] Pieter van Lierop, Erik Lindquist, Shiroma Sathyapala, and Gianluca Franceschini. 2015. Global forest area disturbance from fire, insect pests, diseases and severe weather events. *Forest Ecology and Management* 352 (2015), 78–88.
- [52] Ruihao Wang, Yimeng Liu, and Rolf Müller. 2022. Detection of passageways in natural foliage using biomimetic sonar. *Bioinspiration & Biomimetics* 17, 5 (2022), 056009.
- [53] Albert Weiss. 1990. Leaf wetness: measurements and models. *Remote Sensing Reviews* 5, 1 (1990), 215–224.
- [54] H. Weyl. 1919. Ausbreitung elektromagnetischer Wellen über einem ebenen Leiter. *Annalen der Physik* 365, 21 (1919), 481–500. <https://doi.org/10.1002/andp.19193652104>
- [55] David A. Wikner. 2011. Progress in millimeter-wave imaging. In *RF and Millimeter-Wave Photonics*, Vol. 7936.
- [56] BM Wu, KV Subbarao, AHC van Bruggen, and GGH Pennings. 2001. Validation of weather and leaf wetness forecasts for a lettuce downy mildew warning system. *Canadian journal of plant pathology* 23, 4 (2001), 371–383.
- [57] Muhammet Emin Yanik and Murat Torlak. 2019. Near-Field MIMO-SAR Millimeter-Wave Imaging With Sparsely Sampled Aperture Data. *IEEE Access* 7 (2019), 31801–31819. <https://doi.org/10.1109/ACCESS.2019.2902859>
- [58] Muhammet Emin Yanik, Dan Wang, and Murat Torlak. 2020. Development and demonstration of MIMO-SAR mmWave imaging testbeds. *IEEE Access* 8 (2020), 126019–126038.
- [59] Huaili Zeng, Gen Li, and Tianxing Li. 2024. PyroSense: 3D Posture Reconstruction Using Pyroelectric Infrared Sensing. *Proceedings of ACM UbiCom* (2024).
- [60] Yongzhen Zhang, Yanbo Hui, Ying Zhou, Juanjuan Liu, Ju Gao, Xi-aoliang Wang, Baiwei Wang, Mengqi Xie, and Haonan Hou. 2024. Characterization and Detection Classification of Moldy Corn Kernels Based on X-CT and Deep Learning. *MDPI Applied Sciences* 14, 5 (2024). <https://doi.org/10.3390/app14052166>
- [61] Xiaobing Zhou, Ni-Bin Chang, and Shusun Li. 2009. Applications of SAR interferometry in earth and environmental science research. *MDPI Sensors* 9, 3 (2009), 1876–1912.