UC Riverside

UC Riverside Electronic Theses and Dissertations

Title

Towards Robot-Assisted Precision Irrigation: Proximal Soil Sensing & Physical Leaf Sampling in Orchards

Permalink

https://escholarship.org/uc/item/3p09w857

Author

Campbell, Merrick

Publication Date

2022

Copyright Information

This work is made available under the terms of a Creative Commons Attribution-NonCommercial License, availalbe at https://creativecommons.org/licenses/by-nc/4.0/

Peer reviewed|Thesis/dissertation

eScholarship.org Powered by the California Digital Library
University of California

UNIVERSITY OF CALIFORNIA
RIVERSIDE

Towards Robot-Assisted Precision Irrigation: Proximal Soil Sensing & Physical Leaf Sampling in Orchards

A Thesis submitted in partial satisfaction
of the requirements for the degree of

Master of Science

in

Electrical Engineering

by

Merrick Campbell

June 2022

Thesis Committee:

Dr. Konstantinos Karydis, Chairperson
Dr. Elia Scudiero
Dr. Amit K. Roy-Chowdhury

Copyright by
Merrick Campbell
2022

The Thesis of Merrick Campbell is approved:

Committee Chairperson

University of California, Riverside

Acknowledgments

The author gratefully acknowledges the support of USDA-NIFA under grant # 2021-67022-33453, a UC MRPI Award, and a Frank G. and Janice B. Delfino Agricultural Technology Research Initiative Seed Award. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the funding agencies. Portions of this work have previously appeared in the IEEE 17th International Conference on Automation Science and Engineering (CASE). The author would also like to thank the USDA-ARS U.S. Salinity Laboratory and the University of California Riverside Agricultural Operations for granting access to their fields to conduct experiments.

Thank you:

Dad, for showing me the initial steps from the first Lego bricks through my stubborn teenage years. Mom, for listening through it all. Makayla, for your acts of kindness that brought joy to the journey. Grandpa Chuck and Grandma Billie for sharing the apples from your orchard. I could not have done it without you; I love you all.

John, for giving me my first opportunity to be an engineer and introducing me to agricultural robotics. Nate, thank you for encouraging me to pursue graduate studies. Both of your mentorship will mean more than you will ever know.

Dr. Karydis for challenging me and feeding my ambitions. Amel for keeping it real. Kevin for keeping me on my toes. Hanzhe, Keran, and the rest of ARCS Lab for all the support and late night pancakes.

ABSTRACT OF THE THESIS

Towards Robot-Assisted Precision Irrigation:
Proximal Soil Sensing & Physical Leaf Sampling in Orchards

by

Merrick Campbell

Master of Science, Graduate Program in Electrical Engineering
University of California, Riverside, June 2022
Dr. Konstantinos Karydis, Chairperson

Precision agriculture utilizes sensor networks to inform growers on optimal conditions for applying agronomic inputs (water, fertilizer, pesticides, etc.). Precision irrigation is a subset of precision agriculture that focuses on optimizing water usage. While some growers have embraced these techniques, their usage is far from universal due to cost and labor barriers. Several contemporary works have introduced robotic means for crop monitoring and harvesting. These autonomous systems typically use aerial means for sensing to cover broad regions of crops and ground-based rovers for harvesting. Yet, considerably less work has been performed on using ground-based systems to sample and perform direct measurements. Towards this purpose, two robots are designed and tested. The first robot for proximal soil sensing uses an ECa sensor to generate soil moisture maps of an orchard. The second robot for physical sampling uses a 6-DOF robotic arm to pick a leaf from an avocado tree for stem water potential analysis. These two systems present steps toward a larger robotic system for irrigation measurements in orchard crops.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Precision agriculture techniques allow growers to use a network of sensors to accurately characterize soil-plant-environment processes and monitor current crop conditions [112]. This network can be used to optimize the application of agronomic inputs (water, pesticides, fertilizer, etc.). The quantity of the applied inputs can change dramatically due to variability in environmental conditions (rainfall, temperatures, pests). This variability can have a significant impact on the economic viability and profit margin for growers. Advancing precision agriculture techniques could also mitigate climate change and droughts are impacting farmers. Yet, the adoption of the precision agriculture techniques has been slow despite the advantages. For example, only 12% of growers in the USA use root-zone soil measurements to trigger and budget irrigation events [103, 74].

Precision irrigation is the subset of precision agriculture that focuses on the utilization of water within the farm environment. Field monitoring sensors can increase grower revenue and decrease the environmental footprint of agriculture by applying water directly

where and when required from the ideal source [112, 104]. Several sensors have been developed to both directly and indirectly monitor the water applied to and the moisture content of the fields and orchards. Accurate estimates of water available to plant roots throughout the soil profile can be obtained with soil sensors [96, 14]. While these sensors can accurately determine the available water, they are typically only installed in a few discrete points across many acres. This coverage deficiency fails to capture soil spatial variability, which is a key component of plant-water-environment relationships [69], and leads to an incomplete information on irrigation practices. In addition to discrete localized sensors, remote stand-off sensors can be used to gather data using either unmanned aerial vehicles (UAVs) and satellites [54, 79, 49]. However, these aerial measurements do not fully capture soil moisture information in key regions of interest along the drip-lines and below the surface of the fields. Thus, proximal soil sensing techniques are employed to cover larger areas by moving a sensor over the ground with either a human driven ATV or a robot.

Beyond remote sensing, an increasing number of robots are performing manipulation and interactive tasks with the crop such as harvesting, pruning, and sampling. These robots are often unmanned ground vehicles (UGVs) and focus primarily on harvesting row (e.g., corn and soybean) and orchard crops (e.g., citrus and avocado). Development of agricultural end effectors is an active area of research due to the wide variety of crops and related manipulation tasks. While there are some commonalities across approaches, differences in size, weight, shape, texture, and firmness of specialty crops have led to unique solutions. Some robots can pick strawberries, cucumbers, citrus, and peppers by cutting the stem [35, 102, 101, 6, 7, 88]. These crops can be picked by robots

via wrapping the fruit and twisting it off the stem with either a soft gripper [52, 40, 19], rigid gripper [57, 56, 24, 22, 23, 67, 100], or vacuum [86, 8, 111]. Other robots can pick strawberries, cucumbers, citrus, and peppers by cutting the stem [35, 102, 101, 6, 7, 88].

The use of unmanned ground vehicles (UGVs) for precision irrigation has been a less explored subdomain of agricultural robotics. This thesis explores two related robots for precision irrigation: one for proximal sensing to generate soil moisture maps and another for physical leaf sampling to perform stem water potential analysis. The first robot integrates a apparent electrical conductivity (ECa) sensors on a miniature holonomic platform to profile the salinity composition of the soil (Chapter 3). The second robot uses a custom end effector design (Chapter 4) and a custom perception-actuation framework (Chapter 5) to detect and cleanly cut the stem to remove a leaf from a tree. These robots are preliminary steps towards increasing the sampling frequency and fidelity of precision irrigation data to optimize water usage in orchards.

Chapter 2

Background & Related Works

2.1 Precision Agriculture Techniques

Precision agriculture uses networked sensors to monitor different aspects of the farm environment from soil moisture (SM) content to detecting water leaks. Growers can also leverage this communication network for actuation (i.e. instead of manually turning a valve on a pump, the pump can be triggered via a software application). These networks are usually coupled with GNSS enabled nodes to provide spatial-temporal information about the current farm environment. Terrestrial sensors can also be coupled with satellite data for additional data on crop health [65, 110]. Since Landsat 1 launched in 1972, additional satellites have been launched that contain more sensors and increase data coverage. Today, satellites can provide crop imaging data in green, red, and near-infrared bands (and more) with minimal return times. With the increase in data availability and a decreasing cost of compute resources, the overarching domain of precision agriculture has birthed several subdomains, including precision irrigation (Section 2.1.1). While many specialized sensors

exist for monitoring field health, there are some manual techniques such as stem water potential analysis which would need to be automated for optimal integration within precision irrigation systems (Section 2.1.2).

2.1.1 Precision Irrigation & Soil Sensors

Precision irrigation is the subset of precision agriculture technologies focused on the application of water as an agronomic input. In traditional irrigation practices, water is provided to the crops through surface irrigation (flooding furrows or fields). In more modern irrigation practices, water is deployed using surface-level sprinkler systems, root-zone drip-lines, or subsurface capillary lines [2]. The overarching goal of the precision irrigation systems is to replace open-loop (manual-, volume-, or time-based) control of these irrigation systems with closed-loop control. At a minimum, precision irrigation systems typically involve at least three components: a sensor network to monitor moisture conditions, a decision-making schema, and an actuation system of pumps to disperse water as necessary [34]. Knowing exactly when and how much to water can be challenging due to varying weather conditions and inconsistent soil sensor coverage and measurement frequency.

Soil moisture sensors have several different varieties with unique advantages and drawbacks. Adamchuk et al. provides a comprehensive review of the sensors currently employed in agriculture [3]. Electromagnetic sensors typically measure electrical capacitance, inductance, or conductivity and can be heavily impacted by the tested soil composition. Thus, ground truth calibration should be performed to increase accuracy. (One such sensor is used in the robotic experiments conducted in Section 3.2.1). Radiometric sensors detect the energy reflected or absorbed by soil particles using electromagnetic waves. Like the

electromagnetic sensors, these radiometric sensors can also suffer from calibration issues. While not directly correlated to soil moisture content, mechanical, acoustic, and pneumatic sensors can reveal useful information about the soil composition, compaction, and tillage properties. Electrochemical sensors can also be used to derive information about soil composition.

2.1.2 Stem Water Potential Analysis

While precision irrigation focused on networked sensors, there are still certain data collection techniques for monitoring tree stress levels to inform irrigation routines that are still quite manual. Stem water potential analysis using the Scholander "Pressure bomb" is one such important, yet still largely manual technique [85]. This is a critical process performed by agronomists to estimate tree stress levels and hence optimize irrigation patterns [99]. In this process, an agronomist will cleanly cut the leaf stem and place the leaf inside a pressurized test chamber (colloquially called the "pressure bomb") with its cut end exposed. The agronomist will then gradually increase the pressure until water begins to exit the cut stem [85, 98]. The pressure at which moisture escapes can then be used to determine the leaf water potential and can be correlated with decades of empirical research to optimize irrigation routines [66]. Though effective, the pressure chamber instruments can be tedious and potentially dangerous to operate [76]. Thus, in practice, a single tree is often used to assess orchard health leading to infrequent measurements and undersampled regions.

Though Scholander's pressure chamber has long been the industry standard for measuring stem water potential, contemporary research has tried to develop a less tedious

alternative measurement procedures. Meron et al. utilized laboratory calibrated osmometric sensors embedded in peach and tangerine trees to measure their stem water potential [58]. While their results were comparable to the pressure chamber, the readings were delayed by a few hours and the sensors do require some level of skill to install. Niu et al. developed a measurement device with a low-cost proximate radio frequency sensor combined with machine learning algorithms to track changes in stem water potential [70]. However, the classifications from the network only had an accuracy of 78%. De Bei et al. developed a technique to estimate the water potential of grape leaves using near infrared (NIR) spectroscopy. While his approach could yield rapid results, these sensors require calibration since the structure of the leaf surface can impact the NIR readings. [25]. Methods have also been developed to measure the water potential of a tree by analyzing the behavior and growth of tree trunks [72, 63]. Yet most, if not all, of these methods utilize the pressure chamber as the experimental baseline. Since these sensors are mostly proxies for the water potential, automating the stem water potential analysis via robotic means could lead to higher quality data and higher sampling frequencies than the current manual practice.

2.2 Robotics in Agriculture

The field of robotics as a whole is an interdisciplinary blend of mechanical, electrical, and software engineering focused on autonomous decision-making and precise control of actuated systems. These systems use input devices, such as cameras and other sensors, to interact with an uncertain world. While this field has several unique branches such as path planning, computer vision, and human robot interaction (HRI), agricultural robotics

is a specific subset that contains unique challenges with respect to interacting with dynamic farm environments. Robotics research in precision agriculture has largely focused on remote sensing via ground or aerial robots and harvesting. Section 2.2.1 discusses agricultural robotics focused on monitoring and inspecting crops while Section 2.2.2 discusses some of the work related to manipulating and harvesting produce.

2.2.1 Monitoring & Inspecting

Mobile sensing robots can be used for a wide variety of tasks such as drought stress measurements, pathogen detection, weed detection, nutrient evaluation, crop growth monitoring, and yield prediction [54, 49]. While some of these measurements can be performed with satellites (Section 2.1.1), unmanned aerial systems (UAS) are particularly advantageous since satellites can be hindered by cloud covers and other related weather patterns [82]. Both the satellite and UAS systems can fuse multispectral, hyperspectral, and high-resolution imagery to map crop yield variability [110, 31]. The different classes of UAS platforms, such as fixed-wing and multi-rotor, present tradeoffs in terms of operator ease-of-use, payload capacity, and range [79]. Though fixed-wing aircraft can support a heavier payload over greater distances, they both can support a wide variety of sensors for crop sensing. Unmanned ground vehicles (UGVs) can also support similar sensors and perform crop monitoring tasks. While the UGVs cannot cover terrain as quickly as the UAVs, they can also navigate autonomously and perform high resolution sampling in field environments [16, 95, 97, 94]. These UGVs are often used for probing the fields to generate soil moisture maps [87, 77, 75].

2.2.2 Harvesting & Manipulation

Since harvesting is a high value task for growers, significant research effort has focused on automating this process. Robotic harvesting requires robust methods for both perception (to identify and localize target crops) and actuation (for manipulating and separating the crop from the plant). Such works consider robotic harvesting in both row (e.g., corn and soybean) and tree crops (e.g., citrus and avocado). Despite the unique nature of each type of crop, harvesting end effectors generally have three primary components: the gripping mechanism (mechanical, pneumatic or hybrid), the removal mechanism (mechanical or electrical), and the sensing modality (monocular camera, stereo camera, time-of-flight) [62, 88]. Orchard fruits such as apples and citrus require specific motions to grasp, twist, and pull from the tree without damage [13, 6]. Vine crops such as bell peppers and cucumbers can be directly cut and harvested [51, 7, 102, 101]. More delicate crops like strawberries call for manipulators with force feedback and flexible pneumatic actuators [89, 35, 109, 9]. On the opposite end of the fragility spectrum, end effectors have also been designed for larger and heavier crops such as pumpkins [80].

2.2.3 Leaf Physical Sampling

While most agricultural robotics work has focused on either remote sensing (Section 2.2.1) or harvesting (Section 2.2.2), some work has been performed on physical sampling in the field. Agronomists utilize specialized instruments that can be difficult to transport to the field and thus rely upon sample retrieval for later lab analysis. While this has been mostly a manual process to date, some work has been performed using aerial and ground

robots. Mueller-Sim et al. demonstrated a robotic platform for rapid phenotyping and capable of manipulating leaves for in-situ measurements [64, 1]. Orol et al. developed a tele-operated aerial robot for cutting and collecting leaves from trees [71]. Ahlin et al. presented an algorithm for selecting and grasping tree leaves using a robotic arm [4]. The latter work demonstrates a high level of control using monoscopic depth analysis (MDA) and image-based visual servoing, but focuses on grasping and pulling the leaf instead of cleanly cutting the stem of the leaf, which is useful for stem water potential analysis.

2.3 Computer Vision in Agriculture

Computer Vision is the field of research detected to the signal processing associated with extracting useful information from images and videos. Conceptually, this field is comparable to using a computer to "see" like a human would with their eyes. Broadly, this field has several subdomains including scene reconstruction and understanding, object detection and tracking, and state (pose and motion) estimation. These subdomains traditionally draw from a set of classical techniques rooted in signal processing and image heuristics as well as more modern techniques utilizing neural networks. Computer Vision techniques has been applied alongside precision agriculture and robotics. Section 2.3.1 discusses computer vision for crop health monitoring while Section 2.3.2 discusses computer vision techniques for crop identification, localization, and state estimation.

2.3.1 Crop Health Monitoring

Both classical and machine learning based computer vision techniques have been developed to monitor crop growth, mitigate disease through early detection, and assist with quality control. Zhu et al. developed a computer vision utility to monitor the early growth stages of corn crops. They trained a machine learning classifier to first identify early plant growth and then used scale invariant feature transforms (SIFT) to track the later growth stages [115]. Sadeghi-Tehran et al. used a similar technique to track the growth of wheat. However, they used a bag-of-words approach to separate early growth to distinguish new flowering growth and then a combination of SIFT and a spatial pyramid to extract features for classification with a support vector machine (SVM) [83]. Disease detection is also an active area of research that utilizes computer vision. Jiang et al. utilized a convolution neural network (CNN) to detect five common apple leaf diseases (Alternaria leaf spot, Brown spot, Mosaic, Grey spot, and Rust) [46]. Disease detection is so important to growers, some research has focused on highly parallel algorithms that could be deployed widely on embedded devices such as field programmable gate arrays (FPGAs) [5]. Even after harvesting, computer vision can be used to analyze the crop quality. Su et al. developed a method for determining the quality of potatoes by studying the curvature and volume of over 100 candidate potatoes [92]. Similar produce grading has been performed for tomatoes and other crops [45]. While these techniques are useful within the context of precision agriculture, additional computer vision methods are necessary for interacting with and manipulating the crops.

2.3.2 Crop Identification, Localization, & State Estimation

Manipulating crops requires computer vision systems that precisely and accurately identify and localize target crops within the real-world farm environment [48]. Initial research along these lines identify produce targets by harnessing their distinct colors and shapes. For these approaches, RGB-D cameras that provide both color and depth information can be useful [30, 68, 78]. More recent approaches have focused on utilizing neural network classifiers to identify and track crops from moving video data [17]. Regardless of method, after segmenting the scene and identifying the target crops, they need to be localized spatially within the coordinate system of the robot to be harvested. Thus, determining an estimate of the 6D pose (position and orientation) is critical.

Classical 6D pose estimation techniques traditionally perform feature matching using detected keypoints and then use a RANSAC approach to reconstruct a 3D representation [61, 12]. However, these methods fail when presented with uniformly textured objects and highly occluded scenes. Recent research has explored machine learning methods for generating 3D representations from 2D images [42, 41, 73]. Unfortunately, these learning-based approaches suffer from the same shortcoming: lack of large training datasets. While synthetic data generators can partially mitigate this shortcoming, this approach still requires realistic models with multiple variations in size, shape, orientation or curvature [32, 114].

2.3.3 Leaf Detection & Localization

Leaf detection and localization is a special case of plant segmentation. While most of the research discussed in Section 2.3.2 is motivated by harvest-related manipulation ap-

plications, much of the work associated with leaf detection revolves around path planning for harvesting or leaf segmentation. The process of leaf segmentation can be used to identify leaf components which can be useful for crop health monitoring (see Section 2.3.1). Leaf segmentation can be performed using either classical computer vision tools [18, 60, 27] or machine learning [33, 50, 84]. However, the classical methods are susceptible to environmental changes (changing light, occlusions or hidden surfaces, etc.). The learning-based methods also require large datasets and might not generalize well due to varying environmental factors [36]. These computer vision methods have rarely been deployed on embedded computers as part of a physical sampling system that identifies, localizes, and manipulates a leaf. Such a system presents similar yet unique challenges to crop manipulation, yet presents an inverted problem. Instead of filtering out the leaves to extract the crop, the objective is now to retain the leaves and their individual leaf poses.

Chapter 3

Robotic Proximal Soil Sensing

3.1 Overview

Precision agriculture, and specifically precision irrigation, utilizes near-ground sensing data such as geospatial measurements of soil apparent electrical conductivity (ECa) to optimize water usage. While near-ground sensors provide valuable information to growers, the labor-intensive process of collecting, assessing, and interpreting measurements creates real-world deployment challenges. Performing the measurements with mobile robots could decrese the labor burden while increasing the overall adoption of ECa technology and the accuracy and frequency of data collections. To automate the geospatial ECa measurements and map soil moisture content in micro-irrigated orchard systems, a small wheeled mobile robot was retrofitted with a small electromagnetic induction sensor. The sensor placement considered the potential measurement interference caused by the robot's chassis. The developed software stack integrated the sensor along with a GNSS unit to perform time-stamped geospatial measurements. The prototype robotic ECa measurement method was

evaluated by mapping a 50m x 30m field against a human-conducted measurement baseline obtained by walking the sensor in the same field and following the same path. Despite the robot's small form factor, preliminary experiments suggest that the robotic approach yields comparable measurement to human-conducted ones. This chapter is derived from [15].

3.2 Technical Approach

3.2.1 Measuring Soil Conductivity

Directly measuring soil moisture (SM) is challenging; a reliable proxy for SM spatial variability is geospatial electromagnetic induction (EMI) measurements of soil apparent electrical conductivity (ECa) [20]. These ECa measurements treat the complex interaction of salinity, water content, and soil composition as a bulk conductivity value [20]. The EMI sensor produces a primary electric field that induces eddy currents into moist soil. These eddy currents generate a secondary magnetic field measured by the sensor's receiver to determine the conductivity value [53]. Figure 3.1 depicts the sensor's operating principle. However, these measurements need to be calibrated with a ground-truth baseline to produce accurate SM readings. The ECa measurements can be calibrated at the field-scale with concurrent in-situ measurements of soil moisture using physically-based stochastic modelling to generate accurate root-zone soil moisture maps [69]. Since each orchard contains unique soil compositions, experts can then conduct site-specific ECa-to-SM calibrations using concurrent SM data from either soil cores or discrete, in-situ SM monitoring stations. Once the calibrations have been completed, the soil moisture surveys can be conducted using the

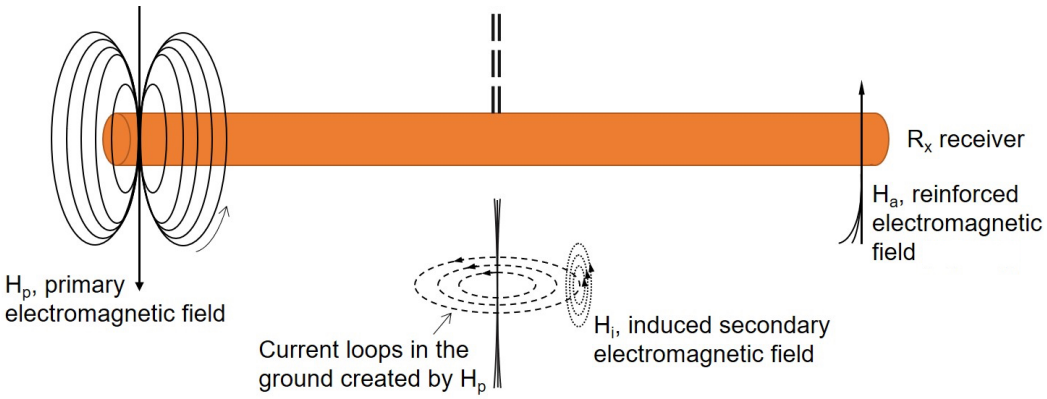

Figure 3.1: Schematic of the operation of the CMD-Tiny probe. Modified from Lesch et al. [53]. This sensor can measure apparent conductivity up to 1000 mS/m with a resolution of 0.1 mS/m with an accuracy of ±4% at 50 mS/m.

EMI sensors.

The SM surveys with EMI sensors are typically conducted manually either by walking the sensor in the field or by driving field vehicles with a trailing sensor (Fig. 3.2). Both approaches are time-consuming, labor-intensive, and limit broad-scale adoption of this technology for frequent SM mapping. Established standard operation procedures [20] recommend that ECa measurements should not be carried out on very dry soils, especially in arid, semi-arid and Mediterranean climates where the space between tree rows of micro-irrigated orchards are generally very dry because irrigation only wets soils very close to the drip emitters [90]. Depending on soil type and irrigation strategy, moist soil is often found up to $1 - 1.5$ m away from drip emitters [91]. In these orchard systems, ECa measurements should thus be carried out close to the trees, along the drip-lines.

The two primary barriers to widespread adoption of ECa surveys using EMI sensors are 1) the cost associated with conducting reliable and consistent surveys [20], and 2) challenges with calibrating the sensor readings to generate an accurate SM estimate. Both

Figure 3.2: Current and envisioned soil moisture (SM) survey techniques using electromagnetic induction (EMI) sensors in traditional and precision agriculture. Current methods include manually-collected data obtained by walking the sensor in the field (left panel), and data obtained by a person driving a field vehicle that pulls the sensor secured on a trailer (central panel). Both methods are labor-intensive. In contrast, the proposed robotic method seeks to automate this process via the use of small and portable agricultural robotics (right panel). The robot prototype is significantly smaller than the ATV and sensor trailer. Due to its small size, it is able to navigate closer to the drip-lines at the base of the trees and exert more control over the spatial component of apparent Soil Electrical Conductivity (ECa) measurements. The robot contains all necessary equipment to perform the measurement including the EMI sensor, a GNSS receiver, and router to provide a local field network.

of these challenges could be mitigated with miniature, portable robotic systems. Robots are ideal for executing precise and repeatable actions such as performing survey measurements. Furthermore, once appropriately calibrated, these robotic systems could automatically re-calibrate to maintain the ability to perform precise and accurate measurements. Prior SM sensing robots encompass very expensive sensor technology, such as cosmic-ray sensors, mounted on large mobile platforms [87, 77, 75]. Such configurations constrain the sensor placement by removing the sensor from the region of interest (close to the trees, along the drip-lines) and reduce the survey frequency due to cost and mobility issues. Thus, miniature, portable, and cost-effective SM survey robots may improve ECa measurement accuracy by bringing the EMI sensor closer to the tree roots and increase survey frequency [44].

3.2.2 Robotic Platform Selection & System Integration

The selection of the base robotic platform for transporting the EMI sensor needs to satisfy several unique design criteria including portability and minimal electromagnetic sensor interference. Since the EMI sensors use an emitter/receiver pair to generate the ECa measurements, they are subject to interference from external magnetic fields. Thus, the placement of the sensor relative to the robot's chassis is a critical design constraint to reduce the possibility that the robot's metal and electrical components could interfere with the ECa measurements. However, the ultimate placement of the sensor relative to the robot is constrained by size, weight, and power (SWaP) requirements to remain portable and effective at traversing uneven terrain. During the traversal, the robot needs to carry a CMD-Tiny (GF Instruments[1]) conductivity sensor which is a small-form sensor representative of many typical sensors on the market. The sensor has diameter 42.5 mm and length 500 mm and weighs 424 g.

A ROSbot 2.0 Pro (Husarion Robotics[2]) [26] wheeled mobile robot was ultimately selected as the base platform due to it's small size, low weight, and high payload capacity. The CMD-Tiny sensor only consumes 4% of ROSbot's total payload capacity of 10 kg. The optimal placement of the sensor relative to the robot was determined empirically to minimize interference on the measurements without sacrificing mobility (Section 3.3.1). ROSbot is also a good candidate for integration due to the system's existing software API. The system's software needs to serve two primary functions: 1) EMI sensor communication and data logging and 2) navigation. For the former, the software also needs to log the recorded

[1]GF Instruments: http://www.gfinstruments.cz
[2]Husarion Robotics: https://husarion.com

ECa soil measurements in conjunction with spatial information such as GNSS. This spatial information is imperative for comparing the measurements with a ground truth baseline as established through core measurements. This data should be stored in a file format such as .CSV that can be processed through several different software tool sets. For the navigation, the robot should be able to be controlled via either manual input or by a waypoint-based trajectory planner.

The system software architecture utilizes ROS as the foundation for automatic geospatial mapping of ECa measurements and soil moisture content in micro-irrigated orchard systems. The system contains ROS nodes for navigation, measurement collection, and data logging. The navigation node utilizes an external GNSS receiver and ROSbot's odometry information fused from motor encoders and IMU to compute the robot's real-time pose (i.e. position and orientation) with an Extended Kalman Filter [47]. The navigation node can receive commands either from manual inputs or a waypoint-based trajectory planner. (Data collected in Section 3.4 used teleoperation.) While not implemented for this analysis, fully autonomous navigation in the field [10, 81, 37, 107] is possible by integrating measurements from the onboard LiDAR [55, 39] and RGB-D camera [28, 43]. A second ROS node extracts SM information from the EMI sensor and publishes the readings. A third ROS node synchronizes the ECa data with GNSS spatial coordinates. This Geospatial SM data is simultaneously displayed in real time and loffed to a local CSV file for post-processing with ArcMap 10.8 (ESRI, Redlands, CA).

3.3 System Validation

3.3.1 Sensor Placement Feasibility Study

The compact and light-weight ROSbot 2.0 Pro was selected as the mobile platform to deploy in the field since it contains the necessary integrated hardware to enable stable maneuverability, remote control, and autonomous navigation. After selecting the platform, the optimal sensor placement needed to be be determined. Figure 3.3 shows how the EMI sensor is orientated with respect to the robot with angle θ, installed with ground clearance d_v, and located away from the robot's back with distance d_h to improve the SM measurement accuracy and reduce the robot's interference. These parameters also constrain the approach angle α and the departure angle β which could limit the capability of overcoming obstacles in the uneven terrain. With careful design parameter selection, the fully-equipped robot can traverse common field terrains with deviation of ± 25 mm. The first step in determining these parameters is to identify a mapping between the distance d_h and EMI interference to determine a viable sensor placement to minimize this interference.

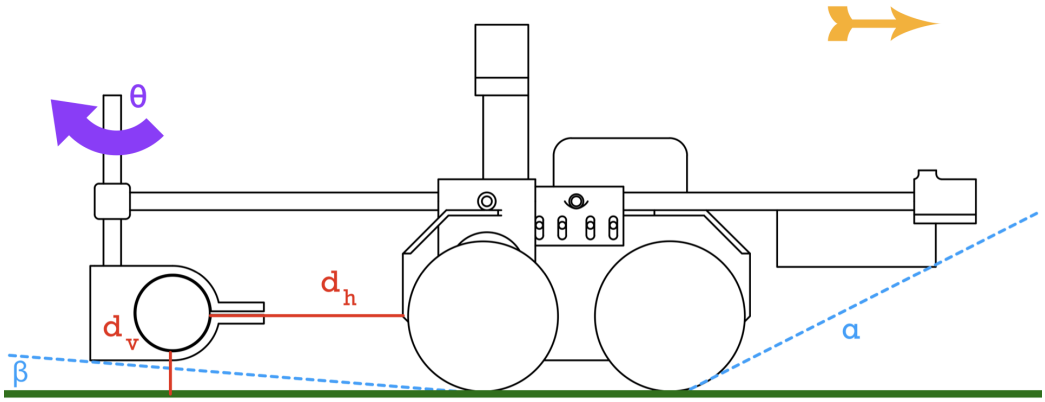

Figure 3.3: The proposed robotic platform. Design parameters d_h, d_v, θ, α and β are determined in this work. The arrow (shown in yellow) denotes the forward-looking direction of the robot.

To determine the ideal sensor placement, the robot was placed at different locations of interest, and at different times and days to establish a rich basis, and several discrete ECa measurements were taken at fixed distance intervals from the robot, $d_h = \{25, 30, 36, 41, 46, 51, 56, 61\}$ cm, and at two orientations, $\theta = \{0°, 90°\}$. Control measurements without the robot present were also taken at all sampled locations. Measurements were taken with two sets of spacers to bound the vertical standoff distance $d_v = \{6.35, 51\}$ mm. These tests were conducted in Riverside, California at the Agricultural Experimental Station of the University of California, Riverside and at the USDA-ARS U.S. Salinity Laboratory with test environments that included grass, bare-soil, and tree roots. These soils were sandy loam with sand content greater than 50%. Future work will measure soil with greater variability in sand, clay, and silt content. Table 3.1 lists all test cases, and Fig. 3.4 depicts an example from testing on irrigated turf. Measurements were performed with the CMD-Tiny sensor in the 'high' mode, which is tuned by the manufacturer

Table 3.1: Conductivity Measurements (mS/m) at Varying Distances d_h (cm)

Test Configuration —— **Distance:**	∞	**25**	**30**	**36**	**41**	**46**	**51**	**56**	**61**
Citrus Grove, $\theta = 0°$	26.7	41.8	35.4	31.9	29.5	28.8	27.9	27.5	27.2
	19.8	30.3	25.9	24.3	22.1	21.1	20.8	20.4	20.1
	21.4	32.7	28.1	26.4	24.3	23.1	22.7	22.2	21.7
	18.3	29.9	28	25.3	23.9	22.5	21.5	21.1	20.2
	18.1	34.3	27.9	23.6	21.2	20	19.3	18.8	18.4
	22.4	33.5	28	26.1	24.7	23.6	22.9	23	22.7
	19.5	35.4	30.1	26.9	23.5	22.5	21.5	21.1	20.6
Irrigated Turf, $\theta = 0°$	30.6	42.7	39.1	35.3	33.5	32.4	31.8	31.3	30.8
	32.6	47.1	42.2	38.2	35.7	34.8	33.7	33.3	32.8
	26.6	41.5	36.1	32.3	29.6	28.8	27.9	27.1	26.9
	28.20	42.90	37.80	33.50	31.60	30.40	29.70	29.20	28.90
	32.20	47.50	41.60	38.70	35.70	34.80	33.90	33.20	33.00
	28.70	39.20	35.10	33.10	31.30	30.40	29.80	29.50	28.90
Irrigated Turf, $\theta = 90°$	28.30	31.20	30.20	29.40	29.10	28.50	28.60	28.40	28.30
	32.40	35.50	33.90	33.20	33.00	32.80	32.80	32.70	32.60
	28.80	31.10	29.70	29.10	28.80	28.60	28.70	28.60	28.40
Tree roots, $\theta = 0°$	14.90	26.30	23.20	20.10	18.30	17.00	16.30	15.80	15.30
	13.9	28.1	23.9	19.6	16.8	15.9	15.2	14.4	14.3
	16.5	31.1	25.2	21.5	19.1	18.3	17.6	17.2	16.9
Tree roots, $\theta = 90°$	14.70	18.30	16.60	15.70	15.20	14.90	14.90	14.70	14.70
Bare Soil, $\theta = 0°$	23.2	39.5	33.5	28.1	26.7	24.9	24.3	23.8	23.5
	22.40	35.00	31.30	28.10	25.90	24.80	23.90	23.40	23.10
	17.20	25.40	23.40	21.20	19.90	19.10	18.50	18.20	17.90
Bare Soil, $\theta = 90°$	22.40	25.20	23.90	23.30	22.90	22.70	22.60	22.50	22.50
	18.50	20.90	20.10	19.60	19.20	19.00	18.90	18.90	18.90

to provide high-accuracy ECa measurements in $[0, 0.7]$ m soil depth.

The tests reveal that the optimal horizontal distance d_h for the sensor would need to be at least 457 mm to minimize the influence of the main robot chassis on the EMI sensor measurements. At smaller distances, the robot affects the sensor measurements but it does not saturate the sensor. Table 3.1 contains the soil conductivity measurements at different test configurations. The column marked with ∞ denotes the control measurement at that location where the robot was not present. Figure 3.5 shows how the ECa measurements differ from the 1:1 control line along with the Pearson coefficient and regression slope and intercept for each distance. As the control vs ROSbot measurements all had a slope very

Figure 3.4: In field tests to identify EMI interference, the sensor (orange cylinder) was operated in a manual configuration to record a series of measurements at multiple distances and fixed orientation (either 0° or 90°) from the robot, over diverse fields. Here we show an instance from testing over irrigated turf at $d_h = 25$ mm and $\theta = 0°$.

close to 1, the influence of the ROSbot on the sensor measurements was considered constant in the range of conductivity measured in this experiment.

After determining d_h, further experiments were conducted to determine the effect of the sensor angle θ. Data suggest that the sensor should be mounted sideways and parallel to the direction of motion of the robot ($\theta = 90°$). While this configuration is ideal from a sensor measurement standpoint, this configuration could make the miniature robotic platform unstable. The final configuration of the sensor on the robot was a compromise between optimal signal to noise ratio and maneuverability. The final prototype constructed in this preliminary research utilized the horizontal sensor orientation with the sensor mounted

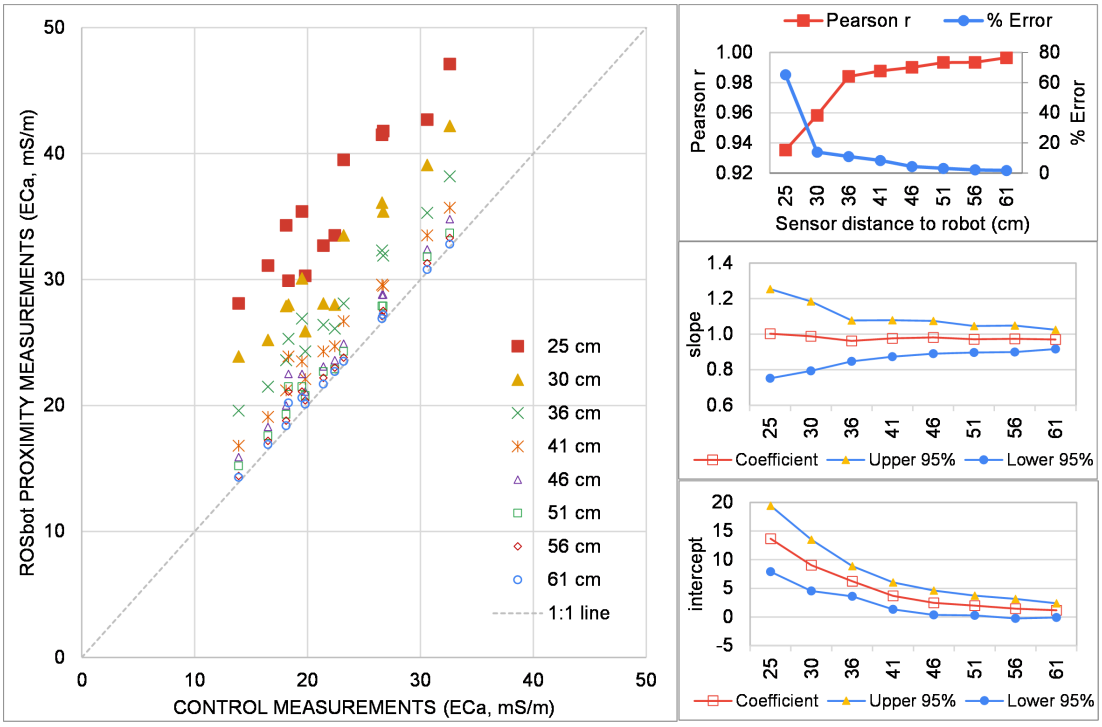

Figure 3.5: For the different field locations, ECa measurements were averaged for each distance d_h. The resulting average values were plotted against the 1:1 control line to determine the constant linear offset. Values that have a line with a slope of 1, parallel to the 1:1 line, and a Pearson correlation close to 1 can be treated as constant linear offsets. Distances greater than 46 cm represent the best candidates for linear offsets, while additional measurements are needed to determine the true relationship at closer distances.

perpendicular to the direction of motion ($\theta = 0°$). Since mounting the sensor at the ideal distance from the robot would reduce mobility, a closer distance of $d_h = 235$ mm was selected and a correction factor was calculated to account for any remaining bias (Section 3.3.1). This correction factor accounts for the small, yet noticeable effect of the robot on the sensor's measurements. Though further refinement of the robotic system could reduce bias, our platform conforms to, and potentially exceeds, standard practices. While the ideal placement would eliminate sensor bias, current best practices often mount the ECa sensors with metal brackets to farm equipment [108], which introduces a measurement bias.

Finally, a set of measurements were conducted to determine the vertical distance d_v by placing the sensor on two spacers: one short spacer (6.35 mm) and another tall spacer (80 mm). The field readings of the ECa values were consistent within the resolution of the measurements between the two spacers. Thus, the value of d_v was determined to be bounded between 6.35 mm and 80 mm. To provide a large margin for the departure angle (as described in Section 3.3.2), the vertical distance d_v was selected to be 50 mm to provide sufficient ground clearance without obstructing the LiDAR's field of view.

3.3.2 Vehicle Integration & Traversal Tests

After the ideal values were determined through field measurements with the EMI sensor and ROSbot, an additional set of validation tests were conducted to verify the platform's mobility. The goal of theses experiments were to determine the maximum practical sensor mounting distance d_h and angles α and β that allow the robot to traverse uneven terrain with deviations of ± 25 mm from level after mounting the sensor without tipping, stalling, or colliding the sensor into the ground. A series of traversal tests were conducted in the lab using foam obstacles on a foam surface to confirm that the robot could indeed navigate the desired terrain (Fig. 3.3.1). Sensor mounting components were designed so that the sensor distance d_h can be adjusted during the traversal tests.

The EMI sensor was mounted on the robotic platform using a combination of carbon fiber rods and custom 3D printed brackets. Using carbon fiber rods is an ideal choice to suspend the sensor behind the robotic platform since the rods have high strength-to-weight ratio, high rigidity, and non-metallic composition. The lightweight nature of

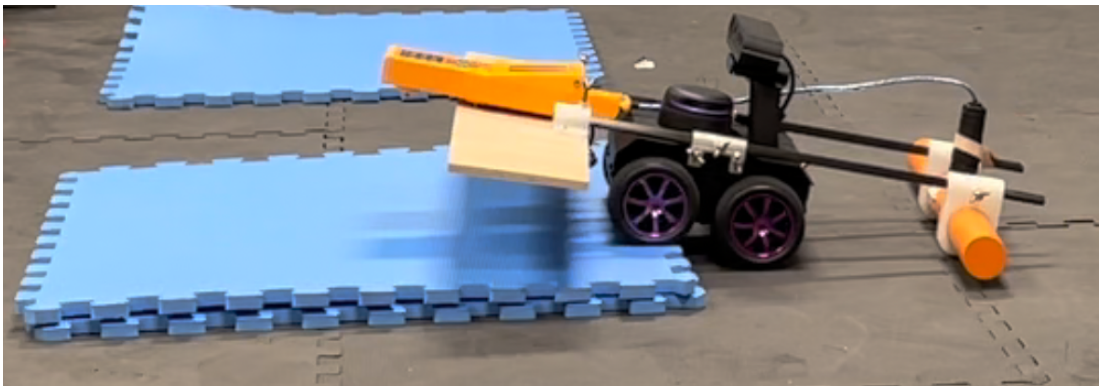

Figure 3.6: To test the system's maneuverability with the sensor attached, an obstacle course was created with several stacks of 12.5 mm thick foam tiles for a total height of 25 mm. The robot was driven over the tiles at various angles and speeds to ensure no unwanted portion of the system made contact with the ground.

carbon fiber aids the design goal of maximum portability. Rigidity is beneficial because this material property minimizes the likelihood that the cantilever sensor arm will vibrate or flex while the robot is traversing obstacles. Finally, metal support rods were not an option since they would affect the conductivity measurements. Use of metal fasteners was minimized to reduce interference with the sensor, though they could not be completely avoided for the sensor mounting.

While the ideal integration configuration would eliminate interference between the robotic platform and the sensor, the sensor placement is constrained by the need to have approach (α) and departure (β) angles that allow the robot to traverse a field with deviations of ± 25 mm from level terrain. As such, the ideal distance d_h determined in Section 3.3.1 is not possible to achieve. Instead, the maximum distance that prevented tipping while still allowing the vehicle to traverse the desired terrain was empirically determined at $d_h = 235$ mm. Following the tests discussed in this and the previous section, and with reference to Fig. 3.3, the selected system mechanical parameters are summarized in Table 3.2.

Table 3.2: Implemented System Mechanical Parameters

d_h	d_v	θ	α	β
235 mm	50 mm	0°	18.4°	12.2°

3.4 Experimental Methods & Results

3.4.1 Establishing a Manual Baseline & Robotic Comparison

Once the preliminary validation tests were complete and the sensor integrated onto the ROSbot platform, the experiments focused on field experiments with the robot. The first experiment aimed to quantify the impact the mounting distance $d_h = 235$ mm has on the ECa measurements and determine if the introduced measurement bias is linear or nonlinear. For this experiment, a row of olive trees was selected as a comparison region between robotic- and manually-collected data. First, the robot with the integrated sensor was driven down the row multiple times to collect ECa measurements along with GNSS coordinates. Then, the same measurements were collected by a human carrying the sensor. These paired measurements were then compared to determine if there are any transients from the robot during operation that could impact the EMI measurement. These tests also served to validate the robot's ability to traverse the field terrain, though these observations were qualitative not quantitative.

Since the mechanical constraints impact sensor placement, the goal here is to determine if the presence of the robot introduced a constant linear offset into the ECa Measurements. Figure 3.7 shows ECa measurements captured via a hand-held manual and robotic data acquisition process. The robotic- and manually-collected datasets were filtered

to remove outliers, averaged, and then compared. The average of the hand-held ECa data set was 16.9 mS/m. The average of the dataset collected with ROSbot was 48.6 mS/m. The average difference was 31.7 mS/m. The slope of the linear relationship between the two datasets was 0.69 with a Pearson correlation coefficient of 0.72. A stronger correlation would be preferred, yet the obtained accuracy offers an acceptable trade-off between sensor measurement accuracy and compact robot design.

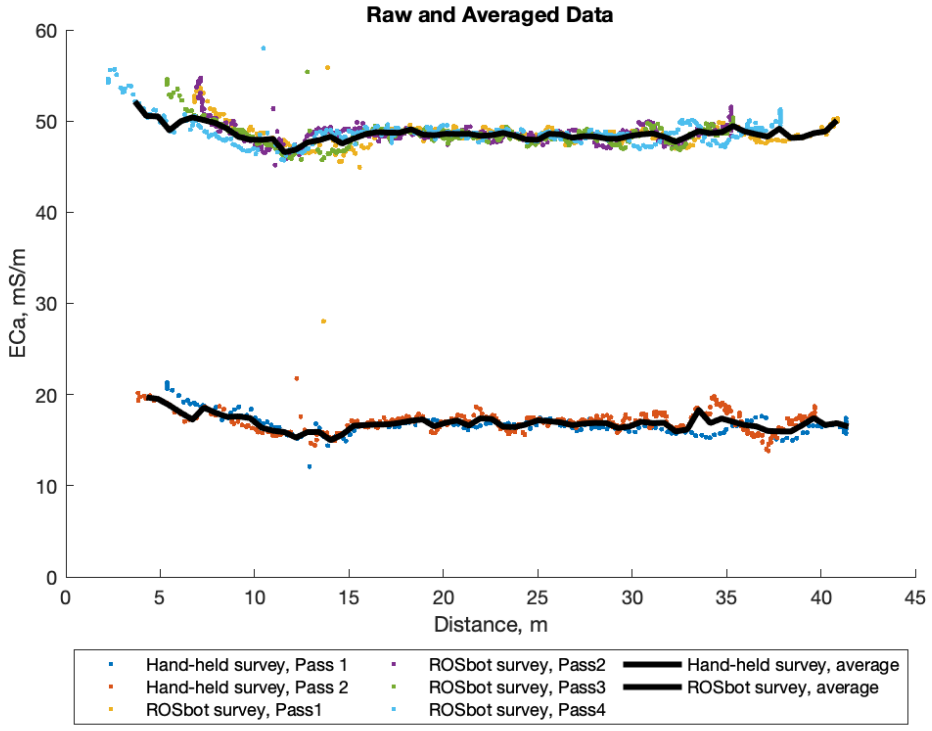

Figure 3.7: ECa Measurements Comparison. Raw data measured with hand-held approach (top) and ROSbot (bottom) are scattered in colors for different trials. Average measurements are plotted in lines.

3.4.2 Generating a Soil Moisture Map

After establishing the measurement baseline, a larger scale survey was conducted to generate ECa maps at a selected study site for comparison of the robotic- and manually-collected measurement data. A study site was selected that encompassed a 50 m $\times$ 30 m subset of a drip-irrigated olive orchard located within the Agricultural Experimental Station of the University of California in Riverside (33°58'24.5"N 117°19'10.3"W). At this location, the soils in the $[0, 0.7]$ m depth range were sandy loam with sand content ranging between 54.6% and 65.9% [11]. Two ECa surveys were carried out on March 12, 2021 at the study site to compare hand-held measurements with these acquired with the integrated robotic soil sensor platform. The GNSS position was recorded along with the ECa measurements to assist in constructing the SM maps for comparison. The robot also recorded its pose (position and orientation), although these measurements were not available during the manual hand-held process. The hand-held survey collected 461 ECa measurements, whereas the ROSbot survey collected 6901. Collected data were sampled at the same frequency, but the robot was moving slower than the human counterpart, thus leading to the increased number of samples via the robotic approach. While the robot collected data at a slower rate than manual operation, this actually provides more accurate measurements [93]. Current commercial practices can suffer from sensor biases and inconsistent practices. A fleet of robots could provide more consistent data and broader coverage of field regions (such as near tree roots, under dense canopies) where a human could not regularly access.

Following collection, the data was processed to generate the SM map from the ECa measurements. These ECa measurements use a bulk soil conductivity model to capture the

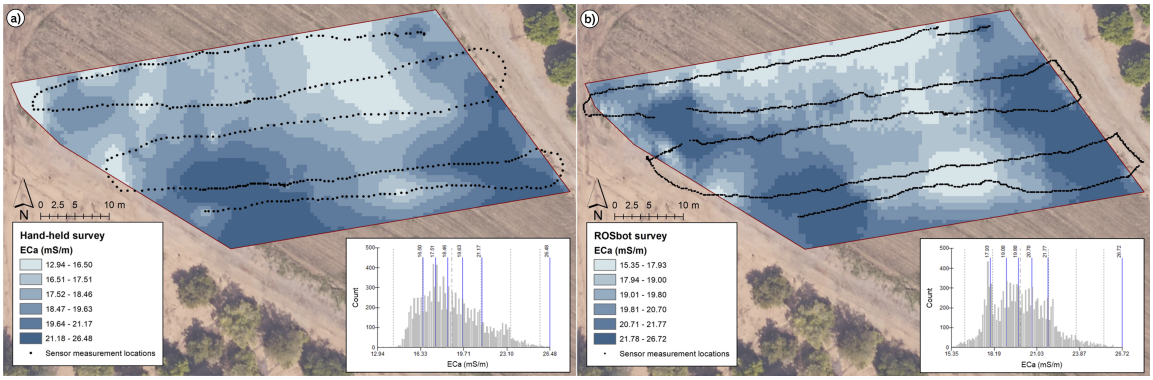

Figure 3.8: Maps of soil apparent electrical conductivity (ECa) for the 0-0.7m soil profile at the study site: a) hand-held survey and b) ROSbot semi-autonomous survey. Color scales are characterized with the quantile method. The maps' ECa frequency distribution are reported on a histogram on the bottom-right of each panel.

complex interplay of soil composition and serve as a good proxy for soil moisture. First, the ECa data were normalized with a natural logarithm transformation following commonly used data filtering procedures [21, 53]. Then, also in line with the standard of practice, any values outside the range of ± 2.5 standard deviation around the average values were deemed as outliers and removed. The hand-held ECa dataset had 2 outliers (0.4%), whereas the ROSbot dataset had 137 (2%). The increased number of outliers in robotic measurements is linked to the fact that those measurements were substantially more compared to manually-collected ones. The average of the hand-held ECa dataset was 19.0 mS/m. The average ECa value measured with the ROSbot was 53.5 mS/m. The difference (34.5 mS/m) was removed from all ECa measurements in the ROSbot dataset to compare the variability and spatial distribution of the two datasets. In ArcMap 10.8 (ESRI, Redlands, CA), the ECa for the two datasets were spatially interpolated using simple kriging with exponential semivariogram to generate maps with 0.5 m $\times$ 0.5 m resolution. These spatial maps are displayed in Fig. 3.8.

The ECa map obtained from the hand-held survey had mean = 18.78 mS/m, standard deviation = 2.31 mS/m, minimum = 12.94 mS/m, and maximum = 26.48 mS/m. The ECa map derived from the ROSbot survey had mean = 19.94 mS/m, standard deviation = 1.85 mS/m, minimum = 15.35 mS/m, and maximum = 26.72 mS/m (Table 3.3). The two maps revealed similar ECa spatial patterns, with the highest ECa values observed at the SE and SW portions of the study site, and the lowest ECa measured in the N and the center of the site. At the pixel-by-pixel level, the maps had a Pearson correlation coefficient = 0.65. This indicates that there are some inconsistencies between the maps, possibly associated with: non-constant influence of the ROSbot on the sensor measurements, different sensor distance to soil and tilt across the two surveys, and higher detail (and variance) captured in the ROSbot survey than in the hand-held one. Observed geolocation inaccuracies are because of the employed low-resolution GPS that also introduces inconsistencies.

Table 3.3: ECa Map Statistics. All values are in mS/m.

Survey	μ	σ	MIN	MAX
Hand-held	18.78	2.31	12.94	26.48
ROSbot	19.94	1.85	15.35	26.72

3.5 Discussion & Outlook

The prototype robot demonstrates the feasibility of conducting ECa surveys using an EMI sensor mounted on a small, portable robotic platform. The ECa geospatial measurements provide real-time spatial information that can be used to infer soil water status. Precision irrigation (i.e. for distinct zones in a field) or traditional irrigation (i.e. the whole

field is managed uniformly) can be scheduled based on soil water status spatial information derived from the prototype platform. The size of the platform is particularly advantageous because it brings the ECa sensor closer to the sample regions of interest such as irrigation drip-lines and tree roots where current handheld, human operated surveys cannot access. The integration approach developed in Section 3.3 highlights some of the design trade-offs with respect to the sensor position relative to the robotic chassis, and quantifies how the robot might bias the sensor readings. Ultimately, the system is capable of gathering the data necessary to create spatial maps and accurate spatio-temporal soil moisture information, which is key to increasing the environmental and economic sustainability of irrigation management in precision and traditional agricultural systems.

Despite the promising initial capabilities, the system could be improved with a few modifications to 1) optimize the relationship between the sensor's signal to noise ratio and platform mobility, and 2) examine sensor performance with different soil types. Additional ECa surveys should sample a wider variety of soil types not present in the initial data set. Current tested soils were sandy loam with sand content greater than 50%, and soil with greater variability in sand, clay, and silt content could produce different calibration conditions. As part of these experiments, future robot iterations should test towing the sensor parallel to robot direction or mounting the sensor in different planes to minimize impact. Although the software foundations are in place with the sensor's ROS integration, the robot's software should implement more intelligent path planning and autonomous navigation capabilities. Finally, robot-gathered ECa data can pair with existing precision agriculture systems to determine optimal irrigation strategies.

Chapter 4

Leaf Cutting End Effector Design

4.1 Overview

The manual nature of stem water potential analysis prevents integration with precision irrigation systems. To conduct this analysis, growers use cleanly cut leaf stems from trees in this industry-standard analysis to determine the irrigation schedules for their orchards. (See Section 2.1.2 for additional details on this methodology.) New end effector designs are necessary for cutting and capturing leaves to enable full automation of stem water potential analysis. While the ultimate goal is to automate this entire process, the first step is to design the mechanism to cut and capture the leaves. While several end effectors exist for harvesting and pruning (Section 2.2.2), no designs currently exist for cutting and capturing leaves for stem water potential analysis. This chapter covers the mechanism design to cleanly cut the stem of a leaf (Section 4.2.1) and retain it, perception component selection (Section 4.2.2) and the mechatronics design to enable safe and repeatable automated stem leaf cutting. The prototype end effector's efficacy is evaluated by cutting leaves

(Section 4.4.2) which are then used for stem water potential analysis with a Scholander pressure chamber, and compared to a manually-cut leaf benchmark (Section 4.4.3).

4.2 Technical Approach

The end effector design is constrained by three key requirements that must be addressed for the leaf water potential analysis. First, the end effector needs to cleanly cut the leaf stem to separate the test specimen from the host tree. Damaging the stem could negatively impact the stem water potential analysis. Second, the end effector needs to capture and retain the cut leaf for analysis. Failing to retain the leaf sample would make analysis impossible. Finally, the end effector needs to maintain a target weight of less than 50% of a typical robotic arm's payload of 2.6 kg to ensure mobility throughout the arm's workspace. While arms can support larger payloads, this payload limit is typical of commercially available collaborative robotic arms.

4.2.1 Cleanly Cutting the Leaf Stem

Stem water potential analysis requires a test leaf specimen with a cleanly cut stem since a damaged or mangled stem would negatively impact the analysis [85]. The force required to cut a stem can be calculated as $F = \pi r^2 \tau$ if the radius of the leaf stem is known. The diameter of ten leaves from four different trees (avocado, clementine, grapefruit, and lemon) were measured to determine a reasonable value for r. The average leaf stem diameter was 2.09 mm with a standard deviation of 0.51 mm. Literature indicates that the shear stress (τ) for cutting organic plant matter ranges from 0.85 to 5.9 MPa [106, 38].

Considering the best and worst case scenarios, upper and lower bounds can be determined for the forces. Using these values, the required force to cut the average leaf ranged from 2.9 to 20 N.

However, organic matter such as leaf stems exhibit visco-elastic properties which means that cutting force needs to be applied at the appropriate rate. When stress is applied to a viscous material, the material resists the deformation linearly with time. When stress is removed from an elastic material, the material returns to the original non-deformed state. Based on visco-elastic material principles, faster cuts will require less force and result in less deformation of the leaf stem. Hence, the rate of cut is equally important to the delivered force. The cutting rate for the leaf stem was determined empirically through a series of experiments (Section 4.4.1). These tests revealed that the minimum cutting speed should be 0.312 m/s. With both a worst-case cutting force estimated and a minimum cutting rate determined, the end effector design target was set to deliver a target force of 20 N at 1.1 m/s. This cutting blade velocity provides sufficient margin over the empirically determined minimum cutting speed of 0.312 m/s to account for any losses and work with a wide variety of tree leaves (e.g., avocado, clementine, grapefruit, and lemon).

4.2.2 Camera Selection and Placement

Several cameras were evaluated to provide a perception component for the end effector (Table 4.1). All the cameras considered could provide both RGB color images as well as depth information. Although the ZED and ZED2 generated high quality images and depth maps, they were excluded because their wide baselines produced fragmented depth maps at close ranges within the tree canopy. The three other cameras were evaluated in

different indoor and outdoor environments. The obtained results show that the Realsense (RS) D435i has the best performance, especially outdoors where it is able to provide a viable depth image at close ranges. Furthermore, this camera produced high-quality point clouds at a lower depth range than the manufacturer specifications (0.1 m). Sample images collected using the RS D435i are shown in Fig. 4.1.

Table 4.1: Candidate Cameras Specifications

Camera	Baseline [mm]	Depth Range [m]	Field of View
ZED	120	0.3 – 25	90° x 60° x 100°
ZED2	120	0.3 – 20	110° x 70° x 120°
ZED mini	63	0.1 – 15	90° x 60° x 100°
RS D435i	50	0.2 – 3	87° x 58° x 95°
RS D455	95	0.4 – 6	87° x 58° x 95°

Two eye-on-hand configurations were considered for the placement of the camera on the end effector (Fig. 4.2). While the configuration in panel (a) can lead to longer look-ahead distances the configuration in panel (b)—angled downward at 45°—was ultimately selected. This configuration strikes a balance between providing useful depth information about the tree (which is needed for obstacle avoidance and navigation around tree branches) and allowing for leaf detection and tracking (which is needed for aligning the end effector with the leaf to cut it). While tree branch avoidance and leaf tracking is not currently implemented, the design herein is cognizant of next steps that need to follow to fully automate leaf stem water potential analysis.

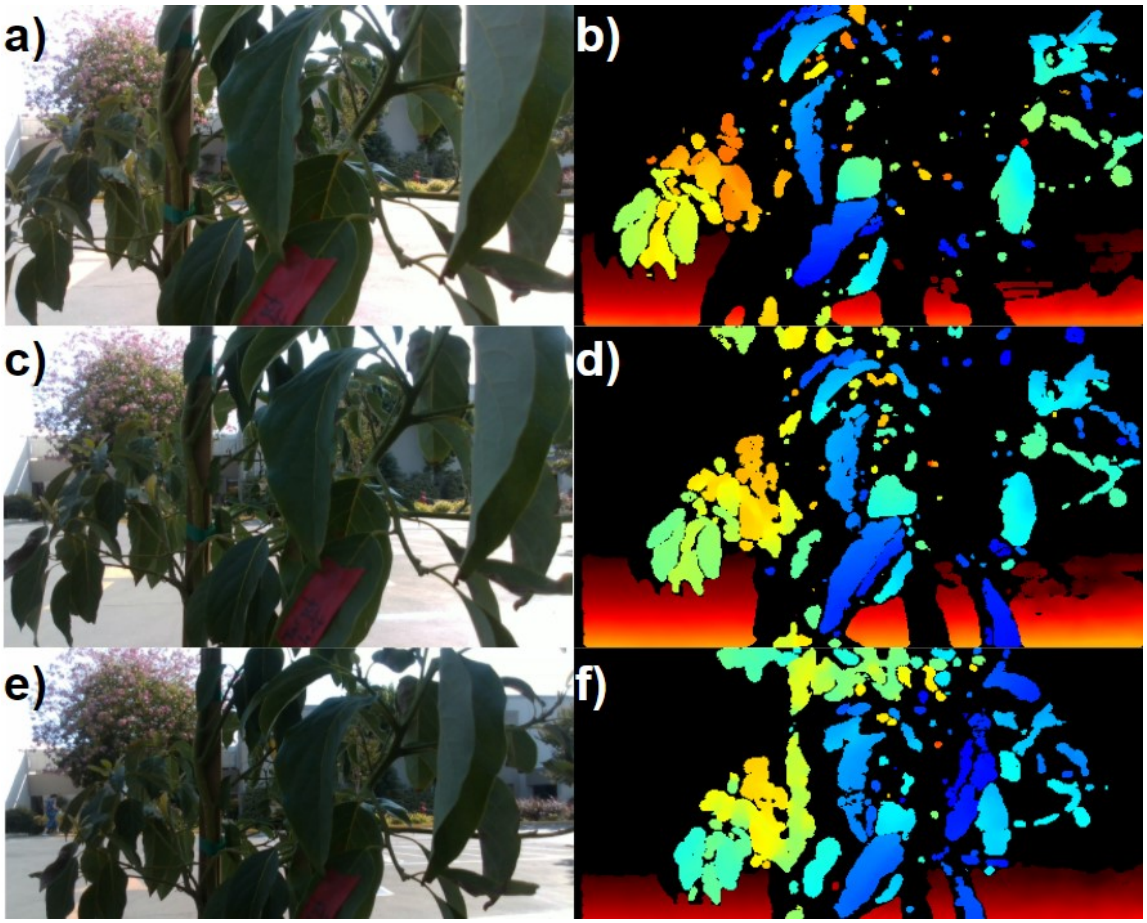

Figure 4.1: Sample RGB and depth images collected from RS D435i in an outdoor environment at (a)–(b) 15 cm, (c)–(d) 20 cm, and (e)–(f) 25 cm.

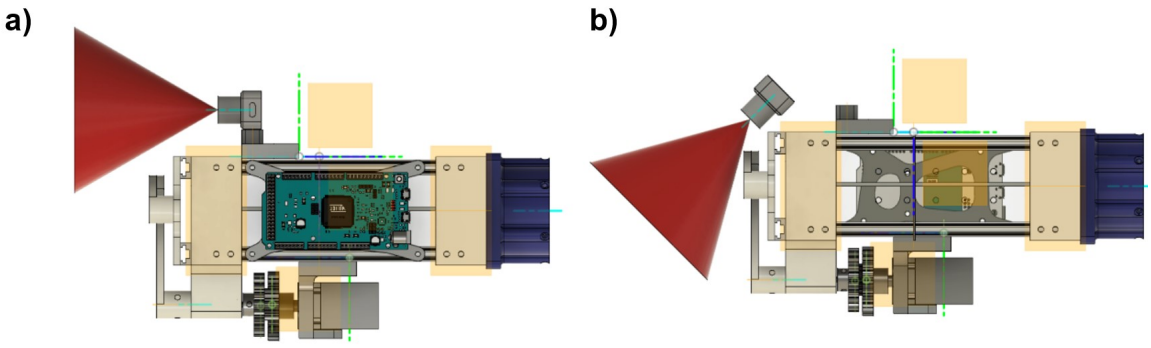

Figure 4.2: Two main camera positions were tested for the end effector: a) straight ahead, and b) angled downward at at 45°.

4.3 System Design

After evaluating the required actuation (Section 4.2.1) and perception (Section 4.2.2) subsystems, an integrated prototype was constructed. The cutting mechanism utilizes two four-bar linkages to actuate a set of sliding gates, one of which contains a razor blade to cleanly sever the stem without damaging the leaf (Fig. 4.3). The gates also help retain the leaf within the end effector's chamber after removal from the tree. These four-bar mechanisms are connected via a geartrain to achieve synchronized motion. A low-cost, high-torque R/C servo (FEETECH FT5335M) drives the geartrain while being amenable to adequate position control. The end effector's chamber has an opening of 110 mm $\times$ 45 mm and a depth of 185 mm to accommodate typical avocado leaves (which are the largest of the four tree crops considered in this chapter). The end effector is constructed with miniature aluminum extrusions, lightweight 3D printed parts, and laser-cut acrylic panels. The assembly weighs a total of 1.09 kg, which is 42% of the robotic arm's 2.6 kg payload.

The end effector operates symbiotically with the Robot Operating System (ROS). High-level control commands are handled via a ROS node running on the robot's embedded computer. This node receives commands from published ROS topics and issues commands to the end effector via Serial UART communication. The end effector contains an embedded microcontroller (Arduino Due) to parse the received serial commands and control the motor that drives the cutting mechanism. A breakout board connected to the Arduino contains a "safe/armed" switch along with LED indicators to reduce the risk of accidental injury from the razor blade. (For redundancy, the high-level ROS control node also has a software "safe/armed" switch.) A 7.4 V 2S LiPo battery powers the end effector mechanism.

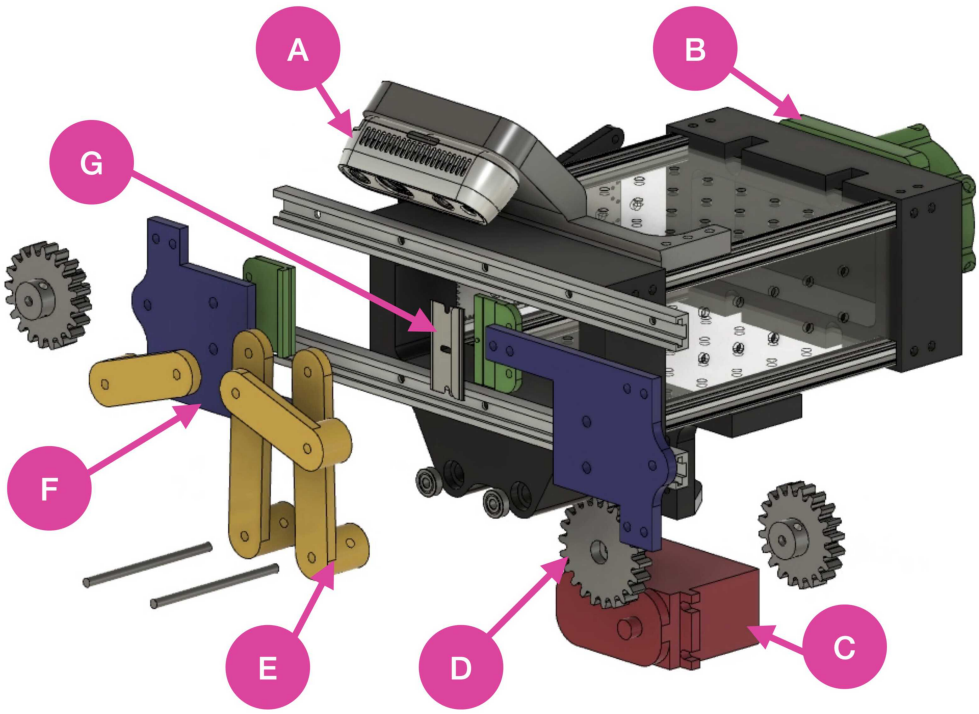

Figure 4.3: Exploded view of the end effector. The design contains several critical components, including the Intel RealSense D435i Depth Camera (A) and an interchangeable robotic arm mount (B). The FEETECH FT5335M R/C servo (C) is connected via a geartrain (D) to four-bar linkages (E). This mechanism closes the gates (F) to cut the leaf with the razor blade (G). This separates the leaf from the tree and retains it within the enclosure for subsequent SWP analysis.

4.4 Experimental Methods & Results

4.4.1 Cutting Speed Tests

To determine the minimum speed necessary to cleanly cut the leaf stem, a preliminary prototype leaf cutter was placed on a level platform above a measuring stick with a high-speed camera positioned to face the cutting blades (Fig. 4.4). Three distinct gear sets were 3D-printed so that they could be inserted between the servo motor and the cutting

mechanism to adjust the speed. For each gear ratio (7:13, 22:13, 41:13), four leaves were inserted into the mechanism and the high-speed camera recorded the cutting attempt for each leaf. The selected motor had sufficient torque margins so that the desired cutting force could be delivered with all tested gearing setups. Recorded frames were analyzed to determine the terminal speed of the cutting mechanism. Since the camera frame rate (240 fps) and the travel distance (19.1 mm) are known, the terminal cutter speed can be calculated as:

$$V = \frac{\delta x}{\delta t} \approx \frac{\Delta x}{\text{frame rate} \times \text{frame count}} \quad . \tag{4.1}$$

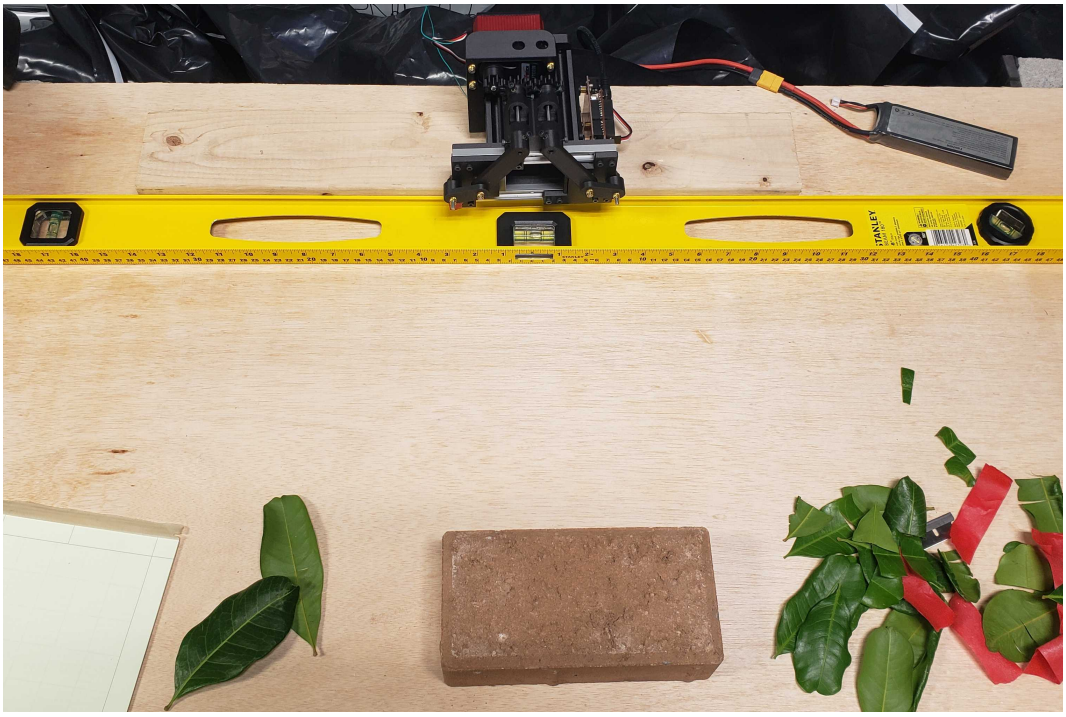

Figure 4.4: Experimental setup for the leaf cutting tests. A first prototype end effector mechanism was placed above a level. A high-speed camera was used to record the cutting operation.

Of the three gear ratios, only the fastest gearing resulted in a cleanly cut leaf. Table 4.2 shows results from all trials and Fig. 4.5 depicts sample frames of the process. From this analysis, we determined that the minimum cutting speed for the cutter should be 0.312 m/s. This result was used in iterative design of the final end effector prototype to optimize the gearing of the four-bar mechanism and the gearing on the servo motor. The final design had a fixed gear ratio that works on the four leaves identified in the preliminary analysis (avocado, clementine, grapefruit, and lemon).

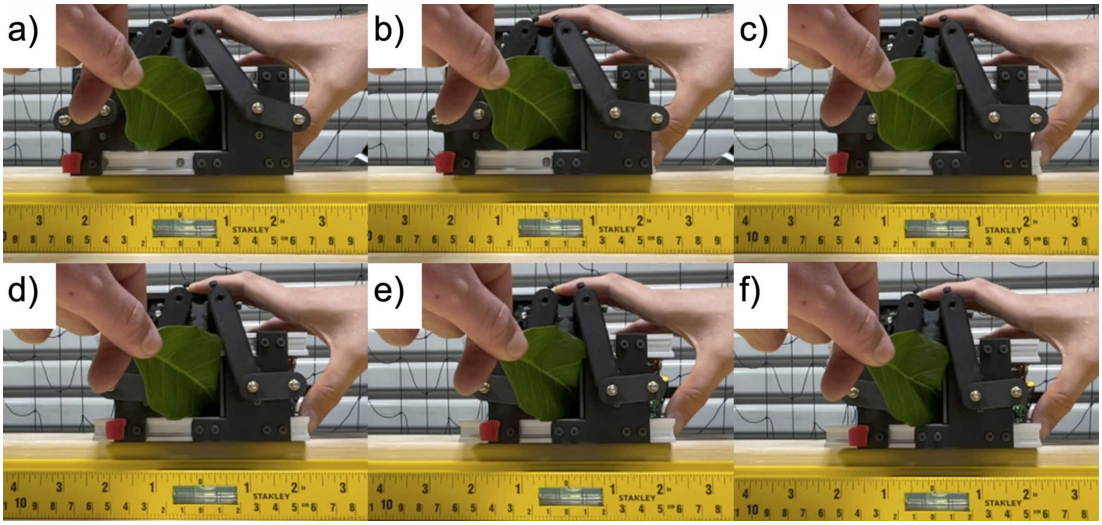

Figure 4.5: Sample frames from one test trial showing the leaf cutting mechanism in action. In frame (f), the mechanism has cut the stem. In these tests, a leaf was already removed from the tree and the mechanism was tested to cut the stem.

4.4.2 Field Leaf Cutting Tests

The integrated prototype developed in Section 4.3 was brought to a local orchard to evaluate the cutting performance on avocado, clementine, grapefruit, and lemon leaves. For these tests, the end effector was manually placed around leaves and activated. Twenty

Table 4.2: Leaf Cutting Velocity Tests

Gear Ratio	Frame Count	Time (s)	Speed (m/s)	Success
7:13	48	0.200	0.095	No
7:13	40	0.167	0.114	No
7:13	39	0.163	0.117	No
7:13	41	0.171	0.112	No
22:13	20	0.083	0.229	No
22:13	25	0.104	0.183	No
22:13	17	0.071	0.269	No
22:13	18	0.075	0.254	No
41:13	16	0.067	0.286	Yes
41:13	20	0.083	0.229	Yes
41:13	11	0.046	0.416	Yes
41:13	14	0.058	0.327	Yes

cutting attempts were performed for each leaf type. For each attempt, a *successful cut* occurs when the enclosed leaf is removed from the tree. A *clean cut* occurs when the leaf is severed cleanly at the stem such that it can be used for stem water potential analysis. The end effector was able to successfully cut 93.75% of the leaves (75 out of 80) with 61.25% being clean cuts. Results are shown in Table 4.3. Lower success rates were observed for the lemon and grapefruit leaves due to their very short stems which make it harder to position the end effector around the stem without interference from branches or other leaves. The end effector worked consistently on clementine and avocado leaves. An instance of one trial and retrieved leaf in the enclosure of the end effector are shown in Fig. 4.6.

4.4.3 Stem Water Potential Analysis Comparison

A series of stem water potential measurements were performed using both manual and end effector cut leaves to verify that the cutting mechanism did not damage the stem.

Figure 4.6: For the field leaf cutting tests, the battery-operated end effector was manually placed about the stem and activated by a button press on the on-board microcontroller (top). Normally, the cut leaf falls into the end effector for retention, but the stem was pulled out for visual inspection (bottom).

Table 4.3: End Effector Field Tests

Crop	Clean cut	Near missed cut	Missed cut
Avocado	15	5	0
Clementine	15	5	0
Grapefruit	11	8	1
Lemon	8	8	4
Rate	61.25%	32.50%	6.25%

Following best practices, all tested leaves were bagged with a water potential reflective foil bag for at least ten minutes before cutting. This helps mitigate transpiration, thus increasing the accuracy of SWP measurements [59]. A total of ten leaves (five for each method) were cut and had their SWP measured.

Preliminary results suggest that the end effector cut leaves produce similar results to the manually cut leaves. The resulting SWP average measurements and standard deviation are tabulated in Table 4.4. The average pressure for the manually-cut leaves was 10.84 Bar, while the end effector cut leaves were 11.34 Bar. The pressure range for the manually-cut leaves is 10.4 to 12.6 Bar, while for the end-effector-cut leaves it ranged from 9.5 to 12.4 Bar. These variations are minor and the manually-cut and end-effector-cut leaves are in essence of same quality for use in SWP analysis.

Table 4.4: Stem Water Potential Measurements

Cut Method	# of Leaves	Avg. Pressure (Bar)	Std. Dev (Bar)
Manual	5	11.34	0.96
end effector	5	10.84	1.14

4.5 Discussion & Outlook

This chapter covered the development of a novel end effector design capable of cleanly cutting the stem of a leaf for use in stem water potential analysis. The system can cut leaves from different types of trees (avocado, clementine, grapefruit, and lemon) with a success rate of 93.75%. However, only 61.25% were clean cuts with suitable stem length for pressure chamber analysis. The enclosure of the end effector allows capturing a bagged leaf without affecting the reliability of the SWP measurement. The average measured pressure for the end effector and manual cuts align closely (cf. 10.84 to 11.34 Bar, respectively). The camera placement on the end effector allows for detection of the leaf on the tree, while providing useful information about the position of the stem.

While the tests conducted in this chapter validate the design of the end effector, the next step is to place the end effector on a robotic arm to determine how the device could be integrated with a leaf picking system. This system could form the foundation of a complete autonomous robotic system to detect, localize, cut, and retain bagged leaves for stem water potential analysis. An optimization of the design will enhance the performance of the system and allow the robot to reach leaves within the canopy closer to the trunk. Chapter 5 seeks to determine the end effector's efficacy when incorporated in such a leaf cutting and retrieval system.

Chapter 5

Action Perception Framework for Leaf Retrieval

5.1 Overview

Picking a leaf requires two key components: actuation and perception. For actuation, a custom-built leaf-cutting end effector (Chapter 4) is mounted on a mobile manipulation base platform (Kinova Gen-2 six degree of freedom [6-DOF] robot arm mounted on a Clearpath Robotics Husky wheeled robot). The perception component utilizes point cloud data from a depth camera (Intel RealSense D435i) for the leaf detection and localization algorithm developed herein (Section 5.2.1). This technique is tested on both indoor and outdoor point clouds from avocado trees. Experimental testing with a real avocado tree demonstrates the proposed approach can enable the mobile manipulator and custom end effector system to successfully detect, localize, and cut leaves (Section 5.3). Figure 5.1

Figure 5.1: The a custom-built end-effector developed in Chapter 4 mounted on an off-the-shelf 6-DOF robotic arm to detect, localize and cut leaves at their stem.

depicts the integrated test setup used for the experiments conducted in this chapter.

5.2 Technical Approach

The perception component revolves around the detection and localization of a viable leaf using RGB-D data. Section 5.2.1 outlines the algorithm and process used to determine a leaf's 6D-pose (position and orientation). This information is then passed to the arm controller to initialize a leaf picking routine (Section 5.2.2). These software systems are run on an Intel i7-10710U CPU. Figure 5.2 highlights how the components interact in a leaf-cutting system.

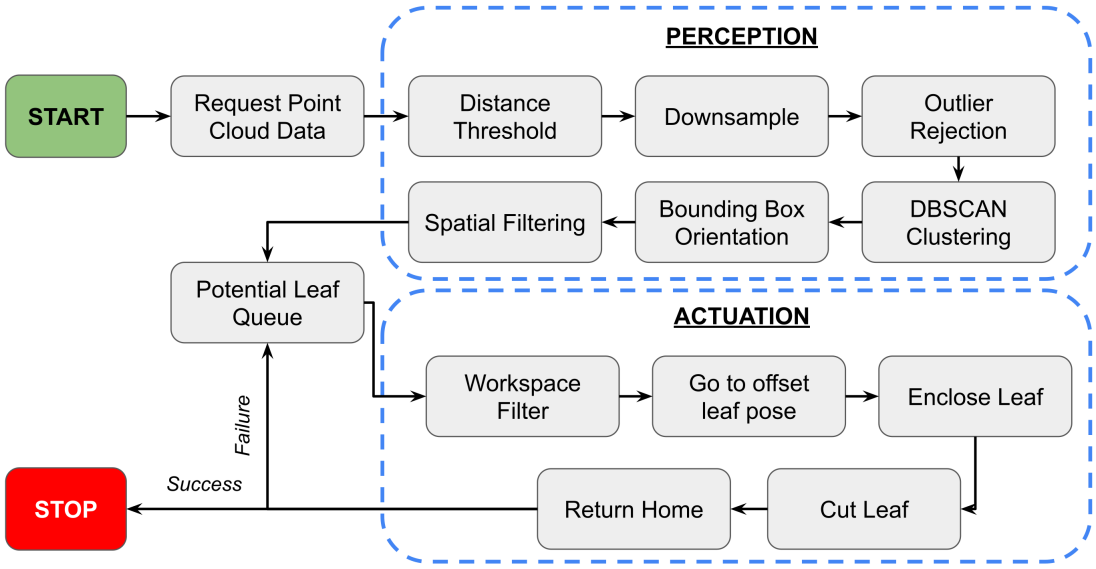

Figure 5.2: The approach jointly considers perception and actuation. The perception module processes point cloud data to segment leaves and deposit leaf candidates into a queue. Candidate leaves are then passed to the robot arm controller to actuate the end effector. If a cut is successful, the routine ends. If unsuccessful, the arm controller requests the next leaf in the queue.

5.2.1 Leaf Detection

The leaf detection and localization pipeline uses a 3D point cloud processed through the Open3D library [113]. The detection phase seeks to obtain the 3D bounding box of leaves candidates from the point cloud captured from the depth camera. First, noisy outliers resulting from sensor measurement inaccuracies are removed. Next, the background is segmented at a specific distance threshold from the camera frame. Then, downsampling is applied to optimize the performance of the upcoming step. Finally, the remaining point cloud segments are grouped into clusters using the Density Based Spatial Clustering of Applications with Noise (DBSCAN) approach [29]. It relies on two parameters, the minimum distance between two points to be considered as neighbors (eps) and the number of mini-

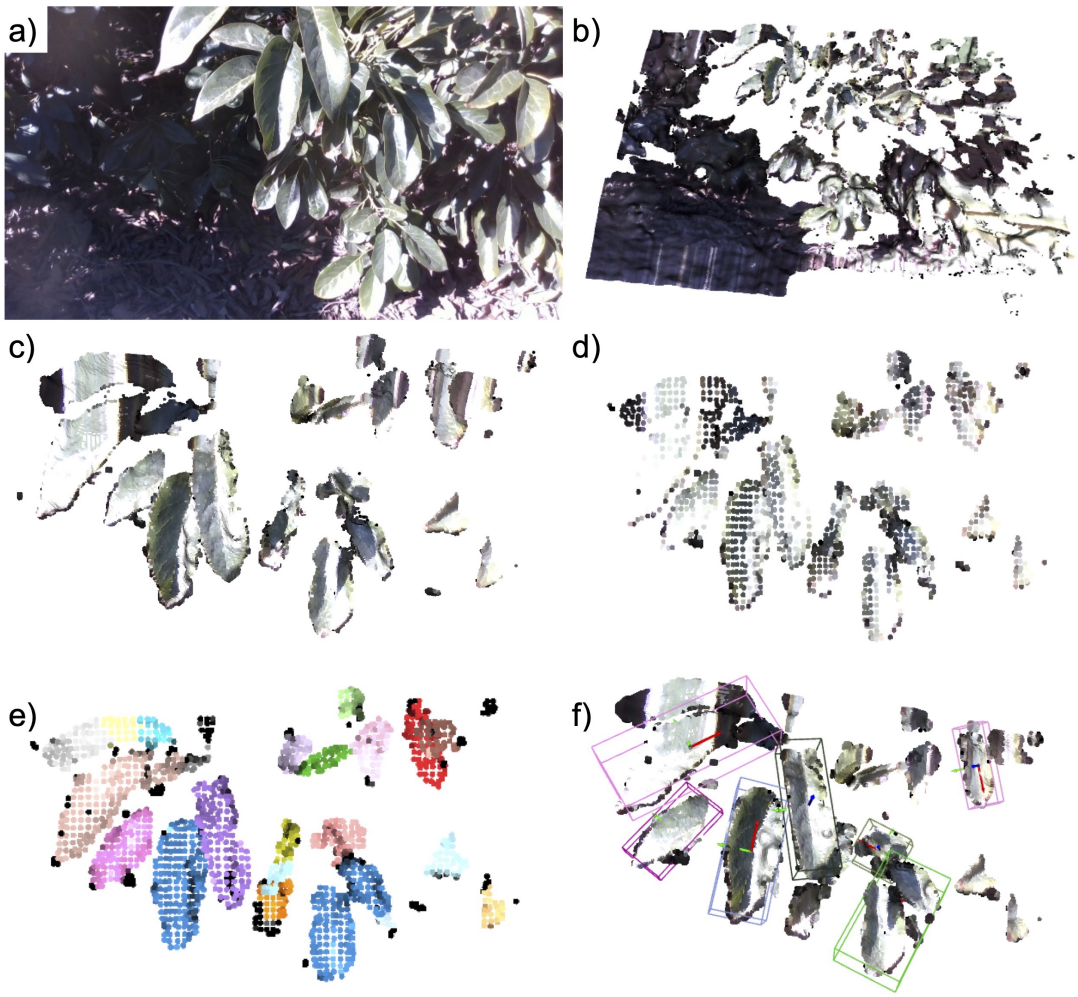

Figure 5.3: Key steps in the proposed leaf detection and localization process. The sample here corresponds to an outdoor point cloud: (a) corresponding RGB image of the tree, (b) raw point cloud, (c) distance filtered ROI, (d) downsampled point cloud, (e) segmented clusters, and (f) detected candidate leaves without 6D pose bounding boxes.

mum points to form a cluster (*MinPoints*). Each resulting cluster is considered a potential leaf and described by a 3D bounding box defined by center $C = [c_x, c_y, c_z]^T$, dimensions $D = [h, w, d]$, and orientation $R(\theta, \Phi, \alpha)$. Then, filtering is applied on the clusters using geometric features of the bounding box: number of points, volume, leaf ratio. Finally, the pose of the center of each bounding box is returned as the 6D pose of a potential leaf.

The detection pipeline was evaluated through offline tests with ROSbags containing indoor and outdoor environments. Indoors (lab with constant light conditions), the Kinova arm collected data with the camera placed at different distances ($0.2 - 0.3$ m) from a potted tree. Outdoors (local orchard with varying light conditions), data was collected manually. A wide range ($0.5 - 1.6$ m) of distances from trees were considered; an example is shown in Fig. 5.3.a. A total of 25 point clouds were collected (10 indoor and 15 outdoor) and tested offline with different combinations for *eps* and *MinPoints* parameters, to determine optimal values for later use. Table 5.1 shows the outcome of the experiments on the 10 indoor point clouds and 15 outdoor point clouds. The pipeline attains an average of 80.0% of detection with a maximum of 90% for indoor dataset, and an average of 79.8% with a maximum 85% for outdoor. Further, the distance between the camera and the tree impacts the optimal values for the point cloud processing. The greater the distance from the camera, the higher *eps* while *MinPoints* decreases.

Table 5.1: Leaf Point Cloud Detection

	Point Clouds	Total # Leaves	Average Detection	Percentage
Indoor	10	20	16	80.0%
Outdoor	15	99	79	79.8%

5.2.2 Arm Control

Identified and segmented leaves serve as target for the arm to move and align the end effector along a viable leaf (to be defined in Section 5.3), at an offset position from the center of the leaf. The offset distance is equivalent to the length of the leaf. Once at

the offset position, the arm moves linearly toward the leaf to capture it. When the leaf is enclosed, the end effector cuts the leaf. Then, the arm returns home.

5.3 Experimental Methods & Results

The overall leaf detection, localization and cutting pipeline was tested with a real potted avocado tree in an indoor laboratory environment. The mobile manipulator and end effector system was positioned at random poses near the base of the tree so that the end effector was at distances ranging between $0.2 - 0.3$ m from the edge of the tree canopy. An experimental trial consisted of collecting a point cloud, storing the identified and localized potential leaves in a queue, and then sending the queued leaves to the arm for a retrieval attempt. Each trial concluded once the queue was depleted and the tree was repositioned for the next trial. Figure 5.4 outlines this process.

For each retrieval attempt, leaf candidates and viable leaves are determined. *Leaf candidates* are leaves that have a pose within the arm's workspace. *Viable leaves* are leaf candidates that have a retrieval path within the arm's workspace. For testing the point cloud detection, both successful captures and successful cuts of the leaf are recorded. A *successful capture* occurs when the end effector is placed around a viable leaf while a *successful cut* occurs when the enclosed leaf is removed from the tree. A *clean cut* occurs when the leaf is severed cleanly at the stem such that it could be used for stem water potential analysis.

Out of 46 trials, 63 potential leaves were detected by the point cloud. (Note that each point cloud in the trial could produce a variable amount of leaves, hence a higher number of potential leaves than trials.) After filtering the potential leaves to remove the

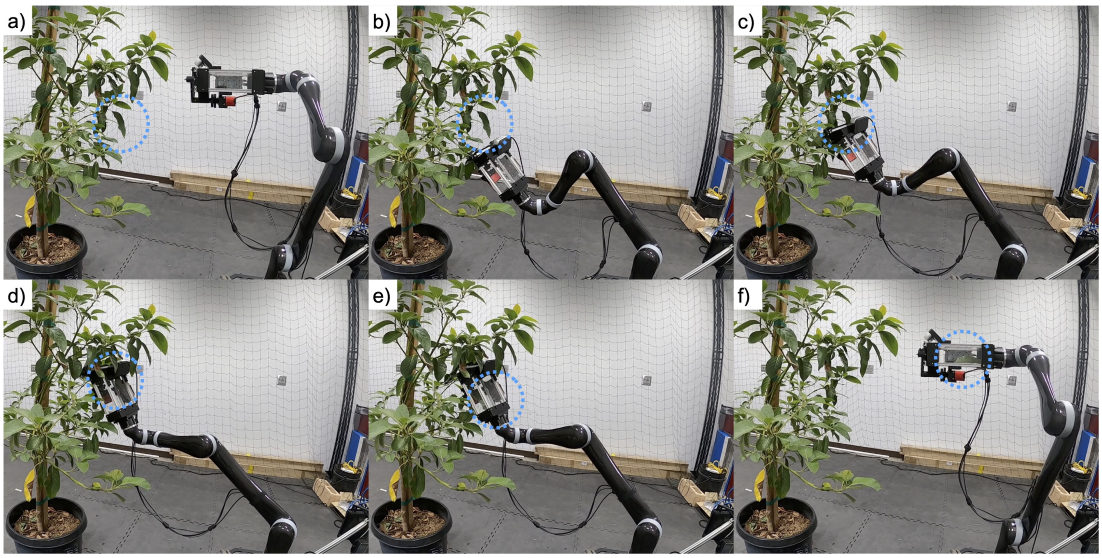

Figure 5.4: Overall leaf retrieval process. During the perception phase, (a) the point cloud is processed to determine a potential leaf. If a viable leaf is detected, (b) the arm will move to an offset position. (c) The arm will then perform a linear motion to capture the leaf. Once in position, (d) the arm will cut the leaf and (e) the leaf will fall into the enclosed chamber. (f) After completing the cut, the arm will return to the home position.

leaves outside of the work space, 39 viable leaves remained. Out of these leaves, 27 were captured successfully (69.2%) while 21 of the 27 captured leaves were cut (77.8%). Table 5.2 summarizes retrieval results while Table 5.3 highlights the process times. The mean point cloud processing (perception) time was 5.6 sec and the mean cutting (actuation) time was 10.6 sec. The mean total retrieval time was 16.2 sec.

Table 5.2: Leaf Retrieval Numbers & Rates

Stage	Number	Rate
Potential Leaves	63	N/A
Candidate Leaves	51	81.0%
Viable Leaves	39	76.5%
Successful Captures	27	69.2%
Successful Cuts	21	77.8%
Clean Cuts	4	19.0%
Near Misses	7	30.0%

The system was able to remove a total of 21 leaves from the tree. However, not all leaves were clean cuts on the stem; four were classified as clean cuts for use in stem water potential analysis. The majority of the leaves were severed at the top of the leaf and not at the stem (Fig. 5.5). The system produced seven near-misses where the stem was leaf was cut within an average of 9.58 mm from the stem (std dev: 6.1 mm). The remaining 10 leaves were severed closer to the middle of the leaf, largely due to collisions with the branches. Similar branch interference also lead to four out of the six missed cuts from the captured leaf. These two problems could be solved in future work through a refined end effector design, more robust path planning to account for branches, and implementing visual servoing for continuous stem alignment as the end effector approaches a viable leaf.

Table 5.3: Leaf Retrieval Performance Time (Seconds)

Metric	Perception Part	Actuation Part	Overall Retrieval
Min	0.5	4.6	6.1
Max	11.0	61.7	62.5
Mean	5.6	10.6	16.2
Median	7.7	8.1	15.3
Std dev	3.9	10.4	10.2

5.4 Discussion & Outlook

This chapter developed a co-designed actuation and perception method for leaf identification, 6D pose estimation and cutting. The 3D point cloud algorithm successfully detected an average of 80.0% of leaves indoors and 79.8% outdoors. Experimental testing of the overall proposed framework for leaf cutting reveals that the system can capture 69.2% of

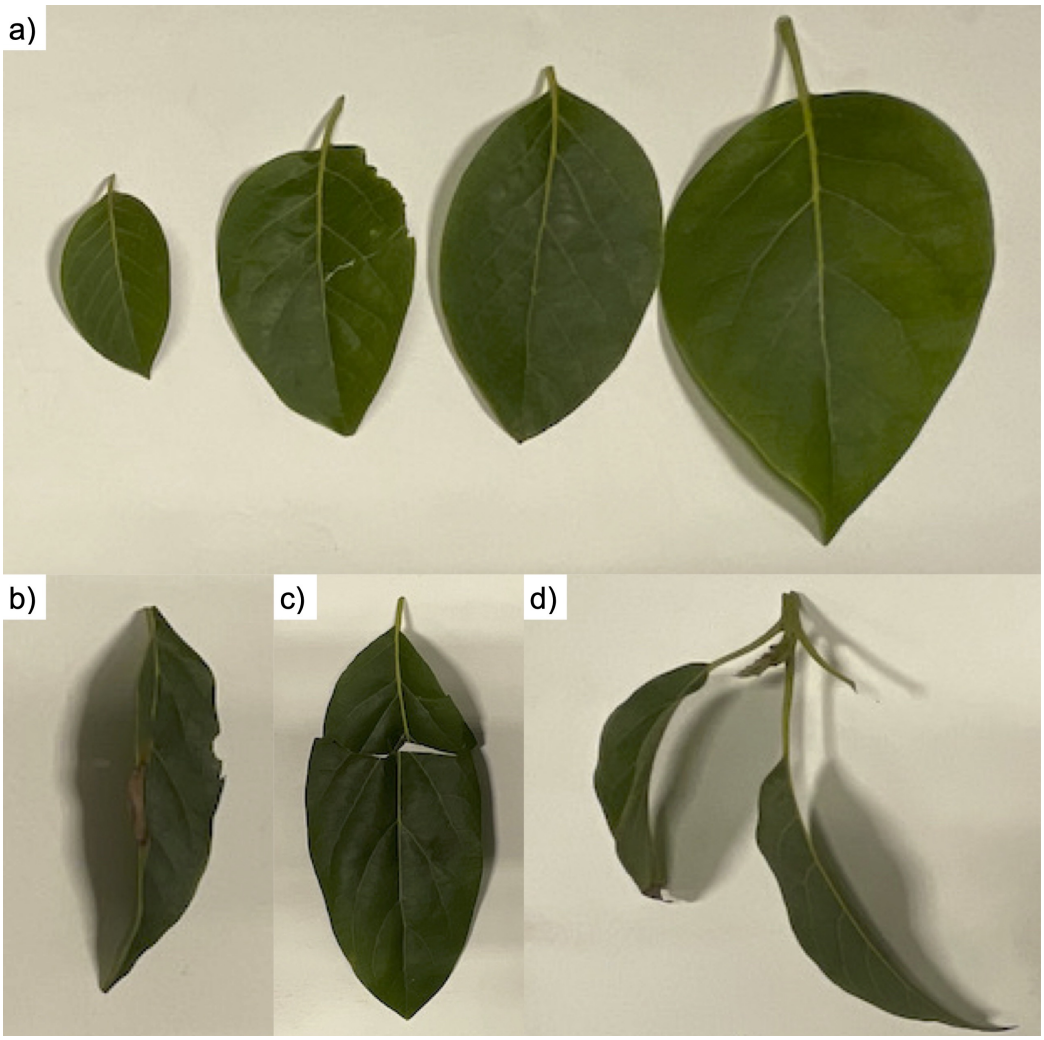

Figure 5.5: Sample leaves cut from the lab's avocado tree during automated indoor tests. (a) The four leaves represent clean cuts suitable for stem water potential analysis. (b) The system also cut seven leaves that were classified as near-misses, which removed the leaf without the stem. (c) The remaining leaves were cut closer to the center, due to interference between the end effector and the branches. (d) In two cases, collateral damage occurred when a second leaf was removed along with the target leaf. These instances were classified as a single successful cut, but not a clean cut since the two leaves would need to be separated for stem water potential analysis.

viable leaves and cut 77.8% of those captured leaves. These results offers a promising initial step toward automated stem water potential analysis, nonetheless several steps remain and are exciting avenues for future work.

While the current path planning approach enables cutting of leaves that are on the periphery of the tree's canopy, leaves closer to the trunk are out of reach. Alternate path planning strategies can be explored to reach leaves within the canopy closer to the trunk by avoiding other leaves, branches, and fruits. Furthermore, implementing visual servoing could better align the cutter with the stem of the leaf as it is about to cut it. These changes would likely require hardware and software co-design of the end effector and integrated robotic system. Finally, to enable automated stem water potential analysis, the captured leaf will need to be transferred from the end effector into a pressure chamber. Once these steps are completed, a robotic leaf sampling system for stem water potential analysis could be fully integrated into a precision irrigation system.

Chapter 6

Conclusions

Robots can increase the frequency and consistency of collecting moisture data for use in orchard precision irrigation systems. Existing irrigation techniques still largely rely upon certain manually collected data which can limit the frequency and consistency of measurements due to labor shortages. The first measurement practice considered for automation was salinity and moisture map generation using an EMI sensor to conduct ECa measurements. The second measurement practice considered for automation was the stem water potential analysis. While other methods such as in-situ and satellite sensors can produce soil moisture information, they each have their benefits and drawbacks. In-situ sensors can provide accurate information, but they rely upon interpolation to provide information across the field. While satellites can also suffer from resolution issues, they can cover large swathes of land frequently. Hence automating near-ground soil moisture data collection and stem water potential analysis could lead to consistent and frequent high-resolution measurements.

Ground-based mobile robots have an important role to play in generating soil moisture maps. A majority of current mapping techniques utilize aerial systems to cover large terrain areas quickly. However, these systems generally fly over the fields and do not conduct measurements close to regions of interest near drip-lines in micro-irrigated orchards. In these types of orchards, measurements are generally performed manually or via an ATV. This manual process leads to inconsistent and infrequent measurements. The design of the mobile robot and experiments conducted in this thesis demonstrate that these measurements can be conducted with a miniature mobile robot. A path for the hardware design has been pioneered, yet significant work remains on fully developing the autonomy and navigation stack to consistently align the robot with the drip lines. While these maps are useful when assessing the soil and field conditions in the orchards, they do not provide information about moisture levels within the trees themselves.

Leaves can be cut from trees for stem water potential analysis to gauge the stress level of the plant. For this analysis to be effective, the stem needs to be cleanly cut and the leaf retained for future analysis. This process has traditionally been performed manually. The first step towards analyzing the process is the design and validation of an end effector to cut and capture these leaves automatically. The end effector developed in this thesis can cleanly cut and capture the leaves. The design was tested in stages to determine the appropriate force and application rate to cleanly cut the visco-elastic leaf stems. These parameters were verified empirically in a laboratory environment before the end effector was taken to an orchard for testing on avocado, clementine, grapefruit, and lemon leaves. While the end effector was able to successfully cut and capture the leaves, the prototype

design is relatively bulky and could be further refined. However, the design enabled the development of an action-perception framework for leaf retrieval.

The leaf retrieval process can be automated using the leaf cutting end effector and a robotic arm. The process developed in this thesis detected and localized leaves on the tree using 3D point cloud data from an RGB-D camera. The related software pipeline produces the 6D pose of leaves within the robotic arm's workspace and then executes the picking commands. This process is able to remove leaves from the tree with varying levels of success. The system was able to successfully capture 69.2% of viable leaves within the robot's workspace by placing the end effector around the target. Of these successful captures, 77.8% were removed from the tree and recorded as successful cuts. However, only 19.0% of these cut leaves would be considered clean enough for stem water potential analysis. Often, branches or other leaves would interfere with the cutting mechanism. These issues could be mitigated by optimizing the end effector size, developing robust path planning for branch avoidance, and implementing visual servoing for stem/cutter alignment. Once these challenges are resolved, a leaf could be picked autonomously by a robot to automate stem water potential analysis and generate useful information for a precision irrigation system.

While the experiments with robotic proximal soil sensing and physical leaf sampling "demonstrate the feasibility" of new robots for precision irrigation systems, significant work remains before field deployment and wide-scale adoption. Challenges remain on both the technical implementation and human acceptance domains. The first technical challenge is the development of fully autonomous control which is the key to offsetting agricultural labor shortages. The second challenge revolves around the data fusion required to integrate

these robots with other sensors within precision irrigation systems. The future of precision irrigation data collection will likely be multi-model and include input from satellites, unmanned aerial vehicles, and ground based mobile robots as explored in this thesis. Even after the technical challenges are resolved, growers may not rush to adopt the robots. New methods can often be met with skepticism and it can be a challenge to replace old habits and routines with new processes. Yet, as the world's demand for food and sustenance increases with the growing population, new techniques will need to be adopted to better allocate diminishing resources. Over time, robots will have a major role to play in optimizing water utilization in farm environments.

Bibliography

[1] J. Abel. In-field robotic leaf grasping and automated crop spectroscopy. Master's thesis, Carnegie Mellon University: Pittsburgh, PA, USA, 2018.

[2] E. A. Abioye, M. Abidin, M. Mahmud, S. Buyamin, M. Ishak, M. Abd Rahman, A. Otuoze, P. Onotu, and M. Ramli. A review on monitoring and advanced control strategies for precision irrigation. *Computers and Electronics in Agriculture*, 173:105441, 2020.

[3] V.I. Adamchuk, J.W. Hummel, M.T. Morgan, and S.K. Upadhyaya. On-the-go soil sensors for precision agriculture. *Computers and Electronics in Agriculture*, 44(1):71–91, 2004.

[4] K. Ahlin, B. Joffe, A.P. Hu, G. McMurray, and N. Sadegh. Autonomous leaf picking using deep learning and visual-servoing. *IFAC-PapersOnLine*, 49:177–183, 2016.

[5] T. Akram, S. R. Naqvi, S. A. Haider, and M. Kamran. Towards real-time crops surveillance for disease classification: exploiting parallelism in computer vision. *Computers & Electrical Engineering*, 59:15–26, 2017.

[6] C. Aloisio, R. K. Mishra, C.Y. Chang, and J. English. Next generation image guided citrus fruit picker. In *IEEE International Conference on Technologies for Practical Robot Applications (TePRA)*, pages 37–41, 2012.

[7] B. Arad, J. Balendonck, R. Barth, O. Ben-Shahar, Y. Edan, T. Hellström, J. Hemming, P. Kurtser, O. Ringdahl, T. Tielen, and B. van Tuijl. Development of a sweet pepper harvesting robot. *Journal of Field Robotics*, 37(6):1027–1039, 2020.

[8] J. Baeten, K. Donné, S. Boedrij, W. Beckers, and E. Claesen. Autonomous fruit picking machine: A robotic apple harvester. In *Field and Service Robotics*, pages 531–539. Springer, 2008.

[9] G. Bao, P. Yao, S. Cai, S. Ying, and Q. Yang. Flexible pneumatic end-effector for agricultural robot: Design & experiment. In *IEEE International Conference on Robotics and Biomimetics (ROBIO)*, pages 2175–2180, 2015.

[10] O. C. Barawid Jr, A. Mizushima, K. Ishii, and N. Noguchi. Development of an autonomous navigation system using a two-dimensional laser scanner in an orchard application. *Biosystems Engineering*, 96(2):139–149, 2007.

[11] D.E. Beaudette and A.T. O'Geen. Soil-web: An online soil survey for california, arizona, and nevada. *Computers & Geosciences*, 35(10):2119–2128, 2009.

[12] E. Brachmann, F. Michel, A. Krull, M. Ying Yang, S. Gumhold, and C. Rother. Uncertainty-driven 6d pose estimation of objects and scenes from a single rgb image. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3364–3372, 2016.

[13] L. Bu, G. Hu, C. Chen, A. Sugirbay, and J. Chen. Experimental and simulation analysis of optimum picking patterns for robotic apple harvesting. *Scientia Horticulturae*, 261:108937, 2020.

[14] G. Calamita, A. Perrone, L. Brocca, B. Onorati, and S. Manfreda. Field test of a multi-frequency electromagnetic induction sensor for soil moisture monitoring in southern italy test sites. *Journal of Hydrology*, 529:316–329, 2015.

[15] M. Campbell, K. Ye, E. Scudiero, and K. Karydis. A portable agricultural robot for continuous apparent soil electrical conductivity measurements to improve irrigation practices. In *IEEE 17th International Conference on Automation Science and Engineering (CASE)*, pages 2228–2234, 2021.

[16] S. Carpin, K. Goldberg, S. Vougioukas, R. Berenstein, and J. Viers. Use of intelligent/autonomous systems in crop irrigation. In *Robotics and automation for improving agriculture*. Burleigh Dodds Science Publishing Limited, 2019.

[17] S. W. Chen, S. S. Shivakumar, S. Dcunha, J. Das, E. Okon, C. Qu, C. J. Taylor, and V. Kumar. Counting apples and oranges with deep learning: A data-driven approach. *IEEE Robotics and Automation Letters*, 2(2):781–788, 2017.

[18] Y. Chen, S. Baireddy, E. Cai, C. Yang, and E. J. Delp. Leaf segmentation by functional modeling. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 2685–2694, 2019.

[19] G. Chowdhary, M. Gazzola, G. Krishnan, C. Soman, and S. Lovell. Soft robotics as an enabling technology for agroforestry practice and research. *Sustainability*, 11(23):6751, 2019.

[20] D. L. Corwin and E. Scudiero. Field-scale apparent soil electrical conductivity. *Soil Science Society of America Journal*, 84(5):1405–1441, 2020.

[21] M. A. Córdoba, C. I. Bruno, J. L. Costa, N. R. Peralta, and M. G. Balzarini. Protocol for multivariate homogeneous zone delineation in precision agriculture. *Biosystems Engineering*, 143:95–107, 2016.

[22] J. R. Davidson, C. J. Hohimer, C. Mo, and M. Karkee. Dual robot coordination for apple harvesting. In *ASABE Annual International Meeting*. American Society of Agricultural and Biological Engineers, 2017.

[23] J. R. Davidson, A. Silwal, C. J. Hohimer, M. Karkee, C. Mo, and Q. Zhang. Proof-of-concept of a robotic apple harvester. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 634–639, 2016.

[24] Z. De-An, L. Jidong, J. Wei, Z. Ying, and C. Yu. Design and control of an apple harvesting robot. *Biosystems Engineering*, 110(2):112–122, 2011.

[25] R. De Bei, D. Cozzolino, W. Sullivan, W. Cynkar, S. Fuentes, R. Dambergs, J. Pech, and S. Tyerman. Non-destructive measurement of grapevine water potential using near infrared spectroscopy. *Australian Journal of Grape and Wine Research*, 17(1):62–71, 2011.

[26] W. de Leeuw, M. Ricke, C. Rosier, T. Wielenga, and S. Grammatico. A multi rosbot laboratory setup for experimenting autonomous driving maneuvers. In *IEEE 28th Mediterranean Conference on Control and Automation (MED)*, pages 532–537, 2020.

[27] B. Elnashef, S. Filin, and R. N. Lati. Tensor-based classification and segmentation of three-dimensional point clouds for organ-level plant phenotyping and growth analysis. *Computers and Electronics in Agriculture*, 156:51–61, 2019.

[28] F. Endres, J. Hess, J. Sturm, D. Cremers, and W. Burgard. 3-d mapping with an rgb-d camera. *IEEE transactions on robotics*, 30(1):177–187, 2013.

[29] M. Ester, H.P. Kriegel, J. Sander, and X. Xu. A density-based algorithm for discovering clusters in large spatial databases with noise. In *International Conference on Knowledge Discovery and Data Mining*, pages 226–231, 1996.

[30] L. Fu, F. Gao, J. Wu, R. Li, M. Karkee, and Q. Zhang. Application of consumer rgb-d cameras for fruit detection and localization in field: A critical review. *Computers and Electronics in Agriculture*, 177:105687, 2020.

[31] C. M. Gevaert, J. Suomalainen, J. Tang, and L. Kooistra. Generation of spectral–temporal response surfaces by combining multispectral satellite and hyperspectral uav imagery for precision agriculture applications. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 8(6):3140–3146, 2015.

[32] M. V. Giuffrida, H. Scharr, and S. A. Tsaftaris. Arigan: Synthetic arabidopsis plants using generative adversarial network. *IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*, pages 2064–2071, 2017.

[33] R. Guo, L. Qu, D. Niu, Z. Li, and J. Yue. Leafmask: Towards greater accuracy on leaf segmentation. In *IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*, pages 1249–1258, 2021.

[34] A. N. Harun, M. Kassim, I. Mat, and S. Ramli. Precision irrigation using wireless sensor network. In *IEEE International Conference on Smart Sensors and Application (ICSSA)*, pages 71–75, 2015.

[35] S. Hayashi, K. Shigematsu, S. Yamamoto, K. Kobayashi, Y. Kohno, J. Kamata, and M. Kurita. Evaluation of a strawberry-harvesting robot in a field test. *Biosystems Engineering*, 105(2):160–171, 2010.

[36] Z. He, W. Feng, X. Zhao, and Y. Lv. 6d pose estimation of objects: Recent technologies and challenges. *Applied Sciences*, 11(1):228, 2021.

[37] M. H. Hebert, C. E. Thorpe, and A. Stentz. *Intelligent unmanned ground vehicles: autonomous navigation research at Carnegie Mellon*. Springer Science & Business Media, 2012.

[38] A. Heidari and G. R. Chegini. Determining the shear strength and picking force of rose flower. *Electronic Journal of Polish Agricultural Universities*, 14(2):13, 2011.

[39] S. A. Hiremath, G. Van Der Heijden, F. K. Van Evert, A. Stein, and C. Ter Braak. Laser range finder model for autonomous navigation of a robot in a maize field using a particle filter. *Computers and Electronics in Agriculture*, 100:41–50, 2014.

[40] C. J. Hohimer, H. Wang, S. Bhusal, J. Miller, C. Mo, and M. Karkee. Design and field evaluation of a robotic apple harvesting system with a 3d-printed soft-robotic end-effector. *Transactions of the ASABE*, 62(2):405–414, 2019.

[41] Y. Hu, P. Fua, W. Wang, and M. Salzmann. Single-stage 6d object pose estimation. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2927–2936, 2020.

[42] Y. Hu, J. Hugonot, P. V. Fua, and M. Salzmann. Segmentation-driven 6d object pose estimation. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3380–3389, 2019.

[43] A. S. Huang, A. Bachrach, P. Henry, M. Krainin, D. Maturana, D. Fox, and N. Roy. Visual odometry and mapping for autonomous flight using an rgb-d camera. In *Robotics Research*, pages 235–252. Springer, 2017.

[44] J. Huang, E. Scudiero, W. Clary, D. L. Corwin, and J. Triantafilis. Time-lapse monitoring of soil water content using electromagnetic conductivity imaging. *Soil Use and Management*, 33(2):191–204, 2017.

[45] G. Jahns, H. M. Nielsen, and W. Paul. Measuring image analysis attributes and modelling fuzzy consumer aspects for tomato quality grading. *Computers and Electronics in Agriculture*, 31(1):17–29, 2001.

[46] P. Jiang, Y. Chen, B. Liu, D. He, and C. Liang. Real-time detection of apple leaf diseases using deep learning approach based on improved convolutional neural networks. *IEEE Access*, 7:59069–59080, 2019.

[47] S. J. Julier and J. K. Uhlmann. New extension of the kalman filter to nonlinear systems. In *Signal processing, sensor fusion, and target recognition VI*, volume 3068, pages 182–193. International Society for Optics and Photonics, 1997.

[48] K. Kapach, E. Barnea, R. Mairon, Y. Edan, and O. Shahar. Computer vision for fruit harvesting robots—state of the art and challenges ahead. *International Journal of Computational Vision and Robotics*, 3:4–34, 2012.

[49] J. Kim, S. Kim, C. Ju, and H. Son. Unmanned aerial vehicles in agriculture: A review of perspective of platform, control, and applications. *IEEE Access*, 7:105100–105115, 2019.

[50] D. Kuznichov, A. Zvirin, Y. Honen, and R. Kimmel. Data augmentation for leaf segmentation and counting tasks in rosette plants. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 2580–2589, 2019.

[51] B. Lee, D. Kam, B. Min, J. Hwa, and S. Oh. A vision servo system for automated harvest of sweet pepper in korean greenhouse environment. *Applied Sciences*, 9(12):2395, 2019.

[52] C. Lehnert, A. English, C. McCool, A. W. Tow, and T. Perez. Autonomous sweet pepper harvesting for protected cropping systems. *IEEE Robotics and Automation Letters*, 2(2):872–879, 2017.

[53] S.M. Lesch, D.L. Corwin, and D.A. Robinson. Apparent soil electrical conductivity mapping as an agricultural management tool in arid zone soils. *Computers and Electronics in Agriculture*, 46(1-3):351–378, 2005.

[54] W. H. Maes and K. Steppe. Perspectives for remote sensing with unmanned aerial vehicles in precision agriculture. *Trends in Plant Science*, 24(2):152–164, 2019.

[55] F. Malavazi, R. Guyonneau, J.-B. Fasquel, S. Lagrange, and F. Mercier. Lidar-only based navigation algorithm for an autonomous agricultural robot. *Computers and electronics in agriculture*, 154:71–79, 2018.

[56] S. S. Mehta, W. MacKunis, and T. F. Burks. Robust visual servo control in the presence of fruit motion for robotic citrus harvesting. *Computers and Electronics in Agriculture*, 123:362–375, 2016.

[57] S.S. Mehta and T.F. Burks. Vision-based control of robotic manipulator for citrus harvesting. *Computers and Electronics in Agriculture*, 102:146–158, 2014.

[58] M. Meron, S. Y. Goldberg, A. Solomon-Halgoa, and G. Ramon. Embedded stem water potential sensor. In *Precision agriculture'15*, page 2. Wageningen Academic Publishers, 2015.

[59] M. Meron, D.W. Grimes, C.J. Phene, and K.R. Davis. Pressure chamber procedures for leaf water potential measurements of cotton. *Irrigation science*, 8(3):215–222, 1987.

[60] T. Miao, C. Zhu, T. Xu, T. Yang, N. L., Y. Zhou, and H. Deng. Automatic stem-leaf segmentation of maize shoots using three-dimensional point cloud. *Computers and Electronics in Agriculture*, 187:106310, 2021.

[61] F. Michel, A. Kirillov, E. Brachmann, A. Krull, S. Gumhold, B. Savchynskyy, and C. Rother. Global hypothesis generation for 6d object pose estimation. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 115–124, 2017.

[62] C.A. Morar, I.A. Doroftei, I. Doroftei, and M.G. Hagan. Robotic applications on agricultural industry. a review. In *IOP Conference Series: Materials Science and Engineering*, volume 997, page 012081, 2020.

[63] A. Moriana, I.F. Girón, M.J. Martín-Palomo, W. Conejero, M.F. Ortuño, A. Torrecillas, and F. Moreno. New approach for olive trees irrigation scheduling using trunk diameter sensors. *Agricultural Water Management*, 97(11):1822–1828, 2010.

[64] T. Mueller-Sim, M. Jenkins, J. Abel, and G. Kantor. The robotanist: A ground-based agricultural robot for high-throughput crop phenotyping. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 3634–3639, 2017.

[65] D. J. Mulla. Twenty five years of remote sensing in precision agriculture: Key advances and remaining knowledge gaps. *Biosystems Engineering*, 114(4):358–371, 2013. Special Issue: Sensing Technologies for Sustainable Agriculture.

[66] A. Naor. Midday stem water potential as a plant water stress indicator for irrigation scheduling in fruit trees. In *III International Symposium on Irrigation of Horticultural Crops 537*, pages 447–454, 1999.

[67] T. T. Nguyen, E. Kayacan, J. De Baedemaeker, and W. Saeys. Task and motion planning for apple harvesting robot. *IFAC Proceedings Volumes*, 46(18):247–252, 2013.

[68] T. T. Nguyen, K. Vandevoorde, N. Wouters, E. Kayacan, J. G. De Baerdemaeker, and W. Saeys. Detection of red and bicoloured apples on tree with an rgb-d camera. *Biosystems Engineering*, 146:33–44, 2016.

[69] D.R. Nielsen, J.W. Biggar, K. T. Erh, et al. Spatial variability of field-measured soil-water properties. *Hilgardia*, 42:215–259, 1973.

[70] H. Niu, T. Zhao, C. Zuber, J. Vasquez-Mendoza, D. Doll, K. Arnold, and Y. Chen. A low-cost stem water potential monitoring method using proximate sensor and scikit-learn classification algorithms. In *ASABE Annual International Virtual Meeting*, page 1. American Society of Agricultural and Biological Engineers, 2020.

[71] D. Orol, J. Das, L. Vacek, I. Orr, M. Paret, C. J. Taylor, and V. Kumar. An aerial phytobiopsy system: Design, evaluation, and lessons learned. In *International Conference on Unmanned Aircraft Systems (ICUAS)*, pages 188–195, 2017.

[72] V. Pagay. Evaluating a novel microtensiometer for continuous trunk water potential measurements in field-grown irrigated grapevines. *Irrigation Science*, 40(1):45–54, 2022.

[73] K. Park, T. Patten, and M. Vincze. Pix2pose: Pixel-wise coordinate regression of objects for 6d pose estimation. *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 7667–7676, 2019.

[74] S. Perdue. 2018 irrigation and water management survey. *Census of Agriculture*, 3, 2017.

[75] A. J. Phillips, N. K. Newlands, S.H.L. Liang, and B. H. Ellert. Integrated sensing of soil moisture at the field-scale: Measuring, modeling and sharing for improved agricultural decision support. *Computers and Electronics in Agriculture*, 107:73–88, 2014.

[76] Pressure chamber safety. `https://www.pmsinstrument.com/maintenance/safety/`. Accessed: 2022-02-28.

[77] J. Pulido Fentanes, A. Badiee, T. Duckett, J. Evans, S. Pearson, and G. Cielniak. Kriging-based robotic exploration for soil moisture mapping using a cosmic-ray sensor. *Journal of Field Robotics*, 37(1):122–136, 2020.

[78] Q. Quan, T. Lanlan, Q. Xiaojun, J. Kai, and F. Qingchun. Selecting candidate regions of clustered tomato fruits under complex greenhouse scenes using rgb-d data. In *International Conference on Control, Automation and Robotics (ICCAR)*, pages 389–393, 2017.

[79] P. Radoglou-Grammatikis, P. Sarigiannidis, T. Lagkas, and I. Moscholios. A compilation of uav applications for precision agriculture. *Computer Networks*, 172:107148, 2020.

[80] A. Roshanianfard and N. Noguchi. Pumpkin harvesting robotic end-effector. *Computers and Electronics in Agriculture*, 174:105503, 2020.

[81] A. Ruckelshausen, P. Biber, M. Dorna, H. Gremmes, R. Klose, A. Linz, F. Rahe, R. Resch, M. Thiel, D. Trautz, et al. Bonirob–an autonomous field robot platform for individual plant phenotyping. *Precision agriculture*, 9(841):1, 2009.

[82] J. D. Rudd, G. T. Roberson, and J. J. Classen. Application of satellite, unmanned aircraft system, and ground-based sensor data for precision agriculture: A review. In *ASABE Annual International Meeting*, page 1. American Society of Agricultural and Biological Engineers, 2017.

[83] P. Sadeghi-Tehran, K. Sabermanesh, N. Virlet, and M. J. Hawkesford. Automated method to determine two critical growth stages of wheat: heading and flowering. *Frontiers in Plant Science*, 8:252, 2017.

[84] H. Scharr, M. Minervini, A. P. French, C. Klukas, D. M. Kramer, X. Liu, I. Luengo, J.M. Pape, G. Polder, D. Vukadinovic, X. Yin, and S. A. Tsaftaris. Leaf segmentation in plant phenotyping: a collation study. *Machine Vision and Applications*, 27:585–606, 2015.

[85] P. F. Scholander, E. D. Bradstreet, E.A. Hemmingsen, and H.T. Hammel. Sap pressure in vascular plants: negative hydrostatic pressure can be measured in plants. *Science*, 148(3668):339–346, 1965.

[86] J. Schupp, T. Baugher, E. Winzeler, M. Schupp, and W. Messner. Preliminary results with a vacuum assisted harvest system for apples. *Fruit Notes*, 76(4):1–5, 2011.

[87] P. Senthil and I. S. Akila. Automated robotic moisture monitoring in agricultural fields. In *International Seminar on Intelligent Technology and Its Applications (ISITIA)*, pages 375–380, 2018.

[88] R. Shamshiri, C. Weltzien, I. A. Hameed, I. Yule, T. Grift, S. K. Balasundram, L. Pitonakova, D. Ahmad, and G. Chowdhary. Research and development in agricultural robotics: A perspective of digital farming. *International Journal of Agricultural & Biological Engineering*, 11(4):1–14, 2018.

[89] W. Simonton. Robotic end effectors for handling greenhouse plant material. *American Society of Agricultural and Biological Engineers*, pages 2615–2621, 1991.

[90] T. Skaggs, T. Trout, J. Simůnek, and P. Shouse. Comparison of hydrus-2d simulations of drip irrigation with experimental observations. *Journal of Irrigation and Drainage Engineering-asce*, 130:304–310, 2004.

[91] T. H. Skaggs, T. J. Trout, and Y. Rothfuss. Drip irrigation water distribution patterns: Effects of emitter rate, pulsing, and antecedent water. *Soil Science Society of America Journal*, 74(6):1886–1896, 2010.

[92] Q. Su, N. Kondo, M. Li, H. Sun, D. F. Al Riza, and H. Habaragamuwa. Potato quality grading based on machine vision and 3d shape analysis. *Computers and Electronics in Agriculture*, 152:261–268, 2018.

[93] K.A. Sudduth, S.T. Drummond, and N.R. Kitchen. Accuracy issues in electromagnetic induction sensing of soil electrical conductivity for precision agriculture. *Computers and Electronics in Agriculture*, 31(3):239–264, 2001.

[94] T. Thayer, S. Vougioukas, K. Goldberg, and S. Carpin. Multi-robot routing algorithms for robots operating in vineyards. In *Proceedings of the IEEE Conference on Automation Science and Engineering*, 2018.

[95] T. Thayer, S. Vougioukas, K. Goldberg, and S. Carpin. Routing algorithms for robot assisted precision irrigation. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 2221–2228, 2018.

[96] G. C. Topp, J. L. Davis, and A. P. Annan. Electromagnetic determination of soil water content: Measurements in coaxial transmission lines. *Water Resources Research*, 16(3):574–582, 1980.

[97] D. Tseng, D. Wang, C. Chen, L. Miller, W. Song, J. Viers, S. Vougioukas, S. Carpin, J. Aparicio Ojea, and K. Goldberg. Towards automating precision irrigation: Deep learning to infer local soil moisture conditions from synthetic aerial agricultural images. In *Proceedings of the IEEE Conference on Automation Science and Engineering (CASE)*, pages 284–291, 2018.

[98] M.T. Tyree and H.T. Hammel. The measurement of the turgor pressure and the water relations of plants by the pressure-bomb technique. *Journal of Experimental Botany*, 23(1):267–282, 1972.

[99] Using the pressure chamber for irrigation management in walnut almond and prune. `https://ucanr.edu/datastoreFiles/391-761.pdf`. Accessed: 2022-02-28.

[100] N. K. Uppalapati, B. Walt, A. Havens, A. Mahdian, G. Chowdhary, and G. Krishnan. A berry picking robot with a hybrid soft-rigid arm: Design and task space control. *Robotics: Science and Systems*, page 95, 2020.

[101] E.J. Van Henten, B.A.J. van Van Tuijl, J. Hemming, J.G. Kornet, J. Bontsema, and E.A. Van Os. Field test of an autonomous cucumber picking robot. *Biosystems Engineering*, 86(3):305–313, 2003.

[102] E.J. Van Henten, D.A. Van't Slot, C.W.J. Hol, and L.G. Van Willigenburg. Optimal manipulator design for a cucumber harvesting robot. *Computers and Electronics in Agriculture*, 65(2):247–257, 2009.

[103] G. Vellidis, V. Liakos, J.H. Andreis, C.D. Perry, W.M. Porter, E.M. Barnes, K.T. Morgan, C. Fraisse, and K.W. Migliaccio. Development and assessment of a smartphone application for irrigation scheduling in cotton. *Computers and Electronics in Agriculture*, 127:249–259, 2016.

[104] G. Vellidis, V. Liakos, W. Porter, M. Tucker, and X. Liang. A dynamic variable rate irrigation control system. In *International Conference on Precision Agriculture*, volume 31, 2016.

[105] G. Vellidis, V. Liakos, W. Porter, M. Tucker, and X. Liang. A dynamic variable rate irrigation control system. In *Proceedings of the 13th International Conference on Precision Agriculture*, volume 13, pages 1–9, 2016.

[106] S. Wang, L. Ren, Y. Liu, Z. Han, and Y. Yang. Mechanical characteristics of typical plant leaves. *Journal of Bionic Engineering*, 7(3):294–300, 2010.

[107] C. Wellington, A. Courville, and A. Stentz. A generative model of terrain for autonomous navigation in vegetation. *The International Journal of Robotics Research*, 25(12):1287–1304, 2006.

[108] C. Willness. Automated soil sensor system adaptable to agricultural equipment, trucks, or all terrain vehicles, Jun 2021.

[109] Y. Xiong, Y. Ge, L. Grimstad, and P. J. From. An autonomous strawberry-harvesting robot: Design, development, integration, and field evaluation. *Journal of Field Robotics*, 37(2):202–224, 2020.

[110] C. Yang, J. H. Everitt, Q. Du, B. Luo, and J. Chanussot. Using high-resolution airborne and satellite imagery to assess crop growth and yield variability for precision agriculture. *Proceedings of the IEEE*, 101(3):582–592, 2013.

[111] K. Zhang, K. Lammers, P. Chu, Z. Li, and R. Lu. System design and control of an apple harvesting robot. *Mechatronics*, 79:102644, 2021.

[112] N. Zhang, M. Wang, and N. Wang. Precision agriculture—a worldwide overview. *Computers and Electronics in Agriculture*, 36(2):113–132, 2002.

[113] Q.-Y. Zhou, J. Park, and V. Koltun. Open3D: A modern library for 3D data processing. *arXiv:1801.09847*, 2018.

[114] Y. Zhu, M. Aoun, M. Krijn, and J. Vanschoren. Data augmentation using conditional generative adversarial networks for leaf counting in arabidopsis plants. In *British Machine Vision Conference (BMVC)*, page 324, 2018.

[115] Y. Zhu, Z. Cao, H. Lu, Y. Li, and Y. Xiao. In-field automatic observation of wheat heading stage using computer vision. *Biosystems Engineering*, 143:28–41, 2016.