

In-Field Robotic Leaf Grasping and Automated Crop Spectroscopy

Justin Abel

CMU-RI-TR-18-51

*Submitted in partial fulfillment of
the requirements for the degree of
Masters of Science in Robotics*

The Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213

August 2018

Copyright © 2018 Justin Abel

Abstract

Agricultural robotics is a growing field of intelligent automation that is proving to drastically increase the speed and reliability of in-field tasks such as precision seed planting, harvesting, field mapping, and crop monitoring. More specifically, plant breeders are beginning to use robotic systems to record the physical traits of crops throughout the growing season at a rate and resolution much higher than the slow and labor intensive manual methods commonly used. Therefore, robotic crop monitoring and automated data collection promise to greatly improve the breeding process by widening this plant phenotyping bottleneck. This thesis describes the development of a custom manipulator capable of grasping leaves for the application of contact-based sensors, as well as the relevant manipulation and computer vision algorithms used for in-field leaf grasping. Furthermore, it details the production of a broadband reflectance spectrometer to capture automated spectral reflectance measurements of leaves using this manipulator. These two systems are mounted to an agricultural, mobile robot platform, developed previously at Carnegie Mellon University, for field testing in sorghum bicolor breeding plots managed by Clemson University throughout South Carolina, USA.

Acknowledgements

I would like to thank Dr. George Kantor, my research advisor, for acting as an incredible mentor, resource, and aid to me over the past two years. Additionally, Dr. Srinivasa Narasimhan and Paloma Sodhi served on my thesis committee and provided valuable feedback on my research work and writing. I would be remiss not to mention Tim Mueller-Sim and Merritt Jenkins for their constant guidance and wisdom throughout my time at CMU. They gave me a harder time than anyone would but also taught me more than anyone could. Lastly, and most importantly, I want to thank my family for constantly providing their love, support, inspiration, and sacrifices so that I can realize my dreams.

Table of Contents

1	Introduction	9
1.1	Motivation	9
1.2	Related Work	10
1.2.1	Hyperspectral Analysis of Crops	10
1.2.2	Automated Harvesting and Manipulation in Agriculture	11
1.2.3	Perception and Leaf Detection	12
1.3	Ground-based Robot Platform	13
1.4	Objectives	15
1.5	Overview	16
2	Design of the Field Reflectance Spectrometer	17
2.1	A Background in Reflectance Spectroscopy and Spectrophotometers	17
2.1.1	Overview of Spectroscopy	17
2.1.2	Overview of Spectrophotometers	18
2.1.3	Light Collection Methods	19
2.1.4	System Requirements	20
2.2	Sensor Selection	21
2.3	Light Source	23
2.3.1	Light Source Background	23
2.3.2	Light Source Selection	25
2.4	System Overview	27
2.4.1	Custom Light Source Design	27
2.4.2	Mechanical and Electrical System Design	29
2.4.3	Software	31
3	Design of the Leaf Grasping Manipulator	32
3.1	The Existing Stalk Grasping Manipulator	32
3.2	Adaptation for Stalk Spectroscopy	33
3.3	Robotic Leaf Grasping and Requirements	34
3.3.1	Problem Statement	34
3.3.2	System Requirements	35
3.4	Custom Manipulator Overview	35
3.4.1	End Effector Design	36
3.4.2	Mechanical Design	38
3.4.3	Electrical Design	39
3.4.4	Software	39

4	Leaf Detection and Manipulation	40
4.1	Leaf Detection	40
4.1.1	Pipeline Overview	40
4.1.2	Basic Stalk Segmentation with Linear RANSAC	42
4.1.3	Detecting Leaves with Parabolic RANSAC	42
4.2	Grasp Point Detection	44
4.2.1	Constraints Due to Limited Degrees of Freedom	44
4.2.2	Graspability Metric	45
4.3	Manipulator Kinematics and Control	48
4.3.1	Inverse Kinematics	48
4.3.2	Manipulator Control	49
4.3.3	Trajectory Generation	50
5	Field Testing and Data Collection	51
5.1	Experimental Leaf Grasping	51
5.2	Spectral Reflectance of Sorghum Stalks	53
5.2.1	Experimental Design	53
5.2.2	Results and Analysis	53
5.3	Spectral Reflectance of Leaves	56
5.3.1	Experimental Design	56
5.3.2	Results and Analysis	56
6	Conclusion	60
	References	61

List of Figures

1.1: Current U.S. renewable energy consumption (in Quadrillion Btu) versus year	9
1.2: Use of custom-made or off-the-shelf components in harvesting projects	11
1.3: Various Crop Harvesting Robots	12
1.4: Rendering of the ground-based agricultural robot developed at CMU	14
1.5: The robotic system during field testing in a bioenergy sorghum breeding plot	15
2.1: Basic construction of a reflectance spectrophotometer	19
2.2: Reflectance spectrophotometer with fiber optic light collection	20
2.3: Black body radiation based on Planck's Law	24
2.4: Thorlabs broadband IR tungsten bulb and spectral output	27
2.5: Custom light source assembly	28
2.6: Light source exploded view	28
2.7: Reflectance spectrometer system	29
2.8: Reflectance spectrometer interior	30
3.1: Stalk grasping manipulator	32
3.2: Updated stalk grasping manipulator with fiber optic attachment	34
3.3: Leaf grasping manipulator	36
3.4: Custom leaf clamp end effector	37
3.5: Fiber optic alignment on the end effector	37
3.6: Leaf clamp actuator design	38
4.1: Example of stereo images and resulting point cloud	40
4.2: Leaf Detection Pipeline	41
4.3: Uncontrollable end effector orientation at various grasp points	44
4.4: Leaf point notation for graspability metric	45
4.5: Visualization of the binary grasp function	46
4.6: End effector vectors plotted over leaf plane normals	47
4.7: Simple planar two link kinematic chain model	48
4.8: Manipulator configurations for multiple grasp points	49
4.9: Four step trajectory for leaf grasping	50
5.1: The leaf grasping manipulator during initial field testing	51
5.2: RVIZ simulation of robotic leaf grasping	52
5.3: Close-up of detected grasp point in RVIZ simulation	52
5.4: Average reflectance spectra of various genotypes	53
5.5: Demonstration of variance in average reflectance spectra	54
5.6: Results of using NDWI and WBI vegetation indices to predict moisture	54
5.7: Correlations from training and test data using neural network for moisture prediction	55
5.8: Correlations from training and test data using neural network for starch prediction	56
5.9: Raw reflectance spectra of leaves from chlorophyll test	57
5.10: Normalized reflectance spectra of leaves from chlorophyll test	58
5.11: Correlations from training and test data using neural network for chlorophyll prediction	58

List of Tables

2.1: Basic System Requirements for Field Reflectance Spectrophotometer	20
2.2: Spectrophotometer Sensors and Specifications	21
2.3: Light Sources and Specifications	26
3.1: Basic System Requirements for Leaf Grasping Manipulator	35
4.1: Key Parameters in the Stalk Detection Algorithm	42
4.2: Key Parameters in the Leaf Detection Algorithm	44

Chapter 1

Introduction

1.1 Motivation

Although agricultural efficiency has continued to improve over the years, the world's population is set to quickly outpace the projected global agricultural output over the next few decades. In fact, the world's population is hypothesized to grow from about 7.5 billion people today to over 9.7 billion by the year 2050, resulting in a projected increase in worldwide agricultural demand of up to 50% [1]. Over the past few decades, the growing pressure to support global food and energy demands has caused significant overexploitation of the world's natural resources resulting in effects such as climate change, deforestation, and land degradation. These unsustainable practices are creating a real concern for food and energy security across the world especially in developing countries [2]. Renewable energy sources offer a much more sustainable alternative to more traditional sources of energy such as coal, natural gas, and other fossil fuels. In the U.S., renewable energy accounts for about 11% of total energy consumption and is growing at an accelerated pace. Furthermore, the use of biomass alone accounts for 45% of total renewable energy consumption [3]. The use of various renewable energy sources from about 1950 to present day can be seen below in Figure 1.1.




Figure 1.1: Current U.S. renewable energy consumption (in Quadrillion Btu) versus year

Source: U.S. Energy Information Administration, *Monthly Energy Review*, Table 10.1, July 2018 [3]

Therefore, it is apparent that the production of bioenergy plays a critical role in continuing sustainable energy production in the country and around the world. However, finding viable sources of bioenergy that can adapt to various climates now and into the future is a challenge. Sorghum, a cereal crop similar to maize, has shown great promise as a bioenergy crop given its genetic diversity, high yield, and tolerance

of marginal conditions [4]. In fact, there are more than 40,000 genetic accessions in the U.S. World Sorghum Collection creating massive potential for the crop to be bred for a wide range of purposes [5].

Traditional forms of plant breeding attempted to produce desirable traits in a plant by creating genetic crosses using simple forms of hand selection and basic knowledge of genetics. The genomic revolution has allowed for fast and inexpensive genome sequencing and gene modification. Modern techniques allow plant geneticists to find markers within the plant genome that correspond to favorable phenotypes to better inform the selection of germplasms for specific breeding applications [6]. However, to closely monitor the association and correlation between genetic markers and physical traits, phenotypic data must be collected on the crop as it grows. Current methods of collecting this phenotypic data include slow, laborious, and often times inaccurate manual methods. This process tends to dramatically slow down the entire breeding process and is thus aptly known as the phenotyping bottleneck [7]. Attempts to alleviate this bottleneck have led to novel high throughput crop phenotyping systems. The goal of these automated systems is to collect data at a much higher rate and resolution compared to traditional manual methods. In order to drive research into optimizing the yield of bioenergy sorghum through intelligent breeding and high throughput phenotyping systems, the U.S. Department of Energy's Advanced Research Project Agency - Energy (ARPA-E) has created the TERRA program [8]. This work detailed by this thesis document and the overall project are funded by the TERRA program to develop intelligent robotic systems used to autonomously collect phenotypic data of bioenergy sorghum.

1.2 Related Work

1.2.1 Hyperspectral Analysis of Crops

Hyperspectral data obtained from crops is frequently used in biological research to make predictions on the biochemical composition and physiological traits of an organic sample. The most common form of hyperspectral data comes in the form of a reflectance spectrum obtained by measuring the intensity of reflected light from a sample as a function of wavelength. Through the use of highly accurate, lab grade spectrographic instruments, high resolution spectra can be recorded. However, these methods are expensive, time consuming, and destructive as the sample must be removed from the field and oftentimes dried and pulverized. These techniques are clearly not suitable for high throughput phenotyping. Instead, portable reflectance spectrometer systems are being used to gather reflectance spectra of crops at a much faster rate and in a non-destructive way. These cheaper, more compact systems come at the cost of accuracy and resolution. Therefore, significant work has been done to develop algorithms that can predict compositional traits that are robust to noise and in-field effects.

Roberts et al. provide a broad overview of state of the art "Hyperspectral Vegetation Indices" [9]. A vegetation index is a number outputted from an algebraic equation that considers reflectance intensity values at one or more wavelengths throughout a reflectance spectrum. These equations are frequently derived from a simple stepwise regression that determines the wavelengths with a high correlation to the trait in question. Therefore, these equations are usually simple ratios, or normalized ratios, such as the equation for NDVI developed by Rouse et al [10] that compares the reflectance at the start of the NIR to the reflectance at visible red light. More specific hyperspectral vegetation index equations have been developed from experiments targeting a specific plant species and thus may not perform well on data set drawn from different species.

Curran et al. provides an early and robust attempt to develop models that predict biochemical compositional traits in leaves [11]. Using the reflectance spectra collected from the leaves of amaranth plants, the group successfully used regression based techniques to determine wavelengths with high correlation to the concentrations of chlorophyll, protein, starch, sugar, amaranthin, and water. The resulting predictive equations resulted in coefficient of determination values above 0.82 between estimated and measured values of chlorophyll, starch, amaranthin, and water. In the same vein, Yendrek et al. collected in-field hyperspectral reflectance measurements, 500 - 2,400 nm, as a means of high-throughput phenotyping to make rapid predictions of key physiological and biochemical traits of maize leaves [32]. Partial least-squares regression techniques were used to develop a predictive measure of chlorophyll, nitrogen, and sucrose content among others that proved to be more accurate than models based on simple spectral indices.

The existence of fully automated, high throughput phenotyping systems that employ hyperspectral sensors are limited but growing. Of particular note is the Field Scanalyzer system designed by LemnaTec consisting of a massive gantry system built around a controlled test plot that carries a payload filled with various sensors throughout the field [12]. These sensors include various cameras, laser scanners, and hyperspectral imagers to name a few. The gantry can position the sensor suite up to about 2 m away from the canopy to provide high resolution images and detailed phenotypic data. A gap does exist in the development of an automated system that can directly contact a crop to acquire hyperspectral data.

1.2.2 Automated Harvesting and Manipulation in Agriculture

The vast majority of manipulation tasks in agricultural robotics deal with automated harvesting. Harvesting robots have received significant attention over the past decade given their enormous potential to cut manual labor costs. However, automating the harvesting process can be quite challenging in practice given variability in crop shape and color, the presence of occlusion and obstacles, as well as unconstrained field effects such as lighting, wind, and rain. Therefore, the design of the manipulator and overall robotic system used for harvesting must be highly specialized for the given application. In a literature review of 50 agricultural robotics papers from 1984 to 2012, Bac et al. found that 78% of manipulators and 98% of end effectors were custom-made [13] as shown in Figure 1.2.

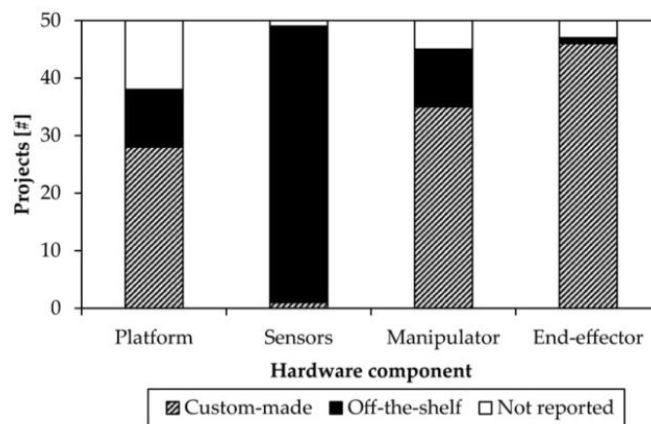



Figure 1.2: Use of custom-made or off-the-shelf components in harvesting projects [13]

Applicable projects that involved custom manipulator design include a novel cherry harvesting system developed by Tanigaki et al consisting of a four degree of freedom manipulator arm that is capable of moving around the full circumference of a tree in order to harvest cherries from all sides [14]. A 3D sensor is mounted onto a second arm to allow for scanning from many viewpoints and thus mitigating the effects of occlusion. A custom manipulator uses suction to position the fruit inside of a two arm gripper that then pulls the cherry from the tree. De-An et al. developed a highly customized, five degree of freedom manipulator optimized for apple harvesting as shown in Figure 1.3a. Apple detection was done through the use of a support vector machine with a radial basis function and manipulator control has aided using visual servoing methods. With the use of a pneumatically actuated spoon-shaped end effector, they were able to achieve a 77% in apple harvesting attempts at an average time of 15 seconds per apple [15].



(a) Apple harvesting robot [15]



(b) Cucumber harvesting robot [16]

Figure 1.3: Various Crop Harvesting Robots

In recent years, however, the ability of off-the-shelf manipulators has improved greatly while the overall mass and form factor has shrunk to a size now suitable for in-field applications. Henten et al. utilized a Mitsubishi RV-E2 manipulator mounted on a custom mobile robotic base, as shown in Figure 1.3b, to harvest cucumbers in greenhouses [16]. A custom end effector with with a padded gripper stabilizes the crop while a thermal cutting device slices the stem in a sanitary fashion. Lehnert et al. created a custom end effector mounted to a Universal Robotics UR5 arm for sweet pepper detection and harvesting [17]. The end effector consisted of a suction cup gripper positioned below an oscillating blade.

The development of agricultural robots capable of crop harvesting is accelerating at a rapid pace. There is no question that the future of farming will include automation performed by these intelligent systems. However, in order to reach a future state where crop harvesting robots are reliable, robust, and affordable, significant work must be done to develop and optimize custom manipulators and end effectors for specialized in-field applications.

1.2.3 Perception and Leaf Detection

Perception in an agricultural setting is challenging for many of the same reasons that make manipulation difficult. These include the wide variability in crop shape and color, heavy occlusion from leaves and other environmental obstacles, fluctuating lighting conditions, and many other unconstrained

field effects. In cucumber harvesting, Henten et al. initially detected cucumbers through the use of a computer vision system that thresholds hyperspectral images at various wavelengths, given the fact that cucumbers exhibit a different reflectance spectra than leaves [16]. Traditional and simpler forms of crop detection and segmentation involving basic color thresholding are more common however in fruit harvesting. However, machine vision in an agricultural setting is not constrained to the task of harvesting. Significant work has been applied to the task of automated yield estimation. Nuske et al. developed computer vision algorithms that are capable of predicting end of season yield in vineyards to within about 10% of actual crop weight [18]. Initial grape detection is done through a radial symmetry transform, and then final grape detection is done using a series of color and texture filters as well as detecting clusters using the KNN algorithm.

In terms of work related to the detection of leaves in an agricultural setting, state of the art research includes the use of geometric model fitting and deep learning, often together, to perform tasks such as leaf segmentation, shape estimation, disease detection, and general phenotyping. Cerutti et al. developed a novel method of leaf segmentation and shape estimation that includes first deriving a set of simple, representative polygonal models then deforming them to detected leaves in an image [19]. Similarly, Xia et al. created a novel technique for detection of multiple leaves in the presence of occlusion and overlapping using active shape models [20]. Leaf segmentation is also performed using region growing techniques on 3D data that classify image points based on surface normals from point clusters as described by Guo et al [21].

Computer vision algorithms have also been applied to high throughput phenotyping of sorghum. Baweja et al. have developed a method to predict stalk count and stalk width using a novel computer vision pipeline consisting of a region based neural network (RCNN) as well as a fully convolutional network (FCN) [22]. These deep learning approaches are becoming more common in crop detection and phenotyping applications over the years. Given that these methods usually require only 2D image data, they are working to simplify the use of complex perception sensors on agricultural robotic systems at the cost of increased computational processing. Lastly, the research of Ahlin et al. provides a fascinating study into leaf detection using convolutional neural networks (CNN) and autonomous leaf picking where manipulator control is advised with visual servoing techniques [23].

The use of computer vision algorithms for perception in agricultural robotics applications has become a necessity. Researchers are developing new techniques ranging from simple color based thresholding for fruit detection to deep neural networks that can estimate anything from stalk width of sorghum to the end of season yield in a vineyard.

1.3 Ground-based Robot Platform

As referenced previously, an agricultural robotic platform was developed recently at Carnegie Mellon University under the TERRA project by Mueller-Sim and Jenkins and described extensively in [24] [25]. The ground-based mobile robot is designed to autonomously navigate rows of sorghum or similar crops planted with a 30 inch row spacing. A rendering of the entire robotic system is shown in Figure 1.4 below.



Figure 1.4: Rendering of the ground-based agricultural robot developed at CMU

The platform carries a series of contact and non-contact sensors for high throughput phenotyping. The base of the robot is about 1.34 m in length, 0.56 m in width, and 1.83 m in height with an overall mass of 140 kg. The system is powered off of a 24 V, 100 Ah LiFePO₄ battery pack that enables the robot to operate for a projected five to eight hours in the field. On board local networks handle the communication between the motor controllers, additional ethernet sensors, and three Intel NUC i7 computers. Software development is done largely with C++ and sensor communication is handled using the Robotic Operating System. On board power regulation is provided by a series of DC-DC converters that take the unregulated battery voltage and output regulated voltages at 24 V, 12 V, and 5 V used by various subsystems.

The robot utilizes a GPS with real time kinematic error corrections for accurate global positioning down to the centimeter level. In addition to the GPS, an inertial measurement unit (IMU) along with dead reckoning estimates from wheel encoders help provide accurate pose estimation during in-field navigation. Autonomous navigation is also aided through the use of a front facing RGB camera as well as a time of flight sensor. These devices are used for live row detection to create a global velocity controller for robust and safe row following. Figure 1.5 below shows the robotic system just outside the entrance to a few rows within a bioenergy sorghum breeding plot in Pendleton, SC.



Figure 1.5: The robotic system during field testing in a bioenergy sorghum breeding plot

A lidar is mounted near the top of the vertical mast that is used to create reconstructions of the sorghum rows as well as crop height predictions. A custom stereo camera is also mounted on the rear of the platform that provides high resolution, 9 MP, images that are used to create accurate estimates of stalk count and stalk width.

Lastly and arguable most relevant to the following research work, a custom three degree of freedom manipulator is mounted to the vertical post of the base. This robotic arm was specifically designed for the task of stalk detection and grasping. One degree of freedom actuates the arm in the vertical direction to grasp stalks at points of various height. The other two degrees of freedom are revolute joints that form a two link, planar kinematic chain to position an end effector to a given point on the plane. This novel agricultural robot is utilized during the field testing of the updated leaf grasping manipulator and reflectance spectrometer systems that are described throughout this thesis.

1.4 Objectives

The main objective of this research is to develop a rapid and robust method of grasping sorghum leaves in the field in order to apply contact-based sensors. A manipulator and suitable end effector capable of accomplishing this task must be selected or designed that is compatible with the pre-existing ground robot platform described above. Furthermore, the potential for using reflectance spectra collected directly from sorghum stalks and leaves to make accurate biochemical and physiological trait predictions is to be explored. Therefore, a compact and portable reflectance spectrometer system to be applied to sorghum must be selection or design. This thesis explores the work associated with accomplishing these objectives.

1.5 Overview

This thesis is organized as follows: Chapter 2 details the mechanical design and sensor selection related to the development of the compact, in-field spectrometer system for high throughput spectral reflectance measurements. Chapter 3 provides a description of the design of a custom manipulator and end effector for the tasks of leaf grasping and direct spectral analysis of leaves. A novel leaf detection pipeline and graspability metric for selecting optimal grasp points are presented in Chapter 4 as well as a brief discussion of the manipulator kinematics and motion planning. Chapter 5 provides an overview of field testing, data collection, and analysis throughout past trips to sorghum breeding plots in South Carolina. Lastly, Chapter 6 gives a summary of the entire thesis and suggests future work.

Chapter 2

Design of the Field Reflectance Spectrometer

2.1 A Background in Reflectance Spectroscopy and Spectrophotometers

2.1.1 Overview of Spectroscopy

At the broadest level, spectroscopy refers to the measurement of electromagnetic radiation due to some physical phenomenon as a function of wavelength or frequency. More specifically, the field of spectroscopy attempts to study the specific interaction between matter and light to try to understand the atomic and molecular structure and chemical composition of various materials. Modern spectroscopy has the ability to analyze properties of light throughout much of the electromagnetic spectrum leading to its prevalence in almost all fields of scientific study.

There is an incredibly wide variety of subfields within spectroscopy that attempt to study the near endless phenomena resulting from the interplay between light and matter. However, the most prominent and arguably impactful fields and forms of spectroscopy are based on a few key interactions between materials and radiation: absorption, emission, and scattering.

Absorption occurs when incident radiation is absorbed by a material and is neither transmitted or reflected out of the material. Radiative energy is absorbed when the frequency of incident photons match the difference in energy between two quantum mechanical states within a material. As a result, a material's absorption spectrum will contain peaks or lines at specific frequencies that correspond to these energy transfer processes. Electronic lines in an absorption spectrum occur when a photon's energy is used to excite a free electron or ion to a higher energy state and are usually found in the ultraviolet and visible light bands. Vibrational lines in a spectrum are caused when incident radiation changes the vibrational state of the molecules and are found in the infrared light band. These vibrations occur when the frequency of radiation matches the fundamental frequency of the bonds between molecules in a material (similar to resonant frequencies found in mass-spring systems). Other forms of energy absorption exist that contribute unique features in a material's absorption spectrum, however it is clear that absorption spectroscopy is a powerful tool that can be used to uncover great information about the composition and structure of a material. It should be noted that absorption and transmission are inherently similar and the spectra of one method are often used to calculate the spectra of the other.

Emission occurs when radiative energy is emitted from a material. Similar to the process of absorption, the energy released by a material is determined by quantum state changes. Therefore, it is the case that a material will absorb and emit radiation at similar and specific frequencies, however, absorption and emission spectra are not equal. Therefore, a spectrum of emitted light can be captured that is unique

to the composition and structure of the material. In practice, this form of spectroscopy requires a way to inject energy into the material before measuring the resulting emission spectrum. This is commonly done through the use of heat, electric stimulation, or radiation.

Scattering occurs when incident radiation is either reflected at the surface of a material or enters through the surface, collides several times with structural components of the material, and is ultimately reflected back out of the material. These interior collisions can be classified into two groups, elastic and inelastic scattering, based on the resulting energy transfer that takes place. Elastic scattering occurs when incident radiation is reflected by the material but the radiation does not lose any energy in the process, thus the photons all remain at their original wavelength and frequency. This process forms the basis of a method of spectroscopy known as reflectance spectroscopy and is also the main focus of the following work. Inelastic scattering occurs when incident radiation is scattered by the material but there is a resulting energy transfer between the matter in the material and the photons. Similar to the process of absorption, inelastic scattering results when an incident photon's frequency matches the resonant frequency of an atom or molecule (where the resonant frequency is based on the energy difference between two quantum states). The odds of a photon of a given frequency striking an atom or molecule with a matching resonant frequency are very low and occur roughly one in ten million collisions. Therefore, the measurement of inelastic scattering is a much more difficult and precise process making reflectance spectroscopy much more common and cheaper to implement.

Reflectance Spectroscopy

As alluded to above, reflectance spectroscopy is the study of reflected or scattered light from a material. When radiation strikes a material, it is either absorbed, transmitted through, or reflected back in the general direction of incidence. Measuring all three of these processes can be incredibly useful and informative techniques in physical and chemical analysis, however, reflectance spectroscopy has a few key advantages that make it particularly fit for in-field automated measurements. Reflectance spectroscopy is nondestructive in the sense that the sample does not need to be altered in any way before taking a measurement. A reflectance spectrometer can simply be placed up to the surface of a material and a reading can be performed, whereas absorption and emission spectroscopy usually require lengthy and destructive material preparation that must be performed in a laboratory setting. Lastly, reflectance spectroscopy is relatively cheap to perform and readings can be done in a matter of seconds. Therefore, reflectance spectroscopy has been chosen as the technique used for crop analysis by the infield robotic system described throughout this paper.

2.1.2 Overview of Spectrophotometers

Similar to the name of the field, a spectrometer is a tool used to measure the spectral response of a physical or chemical phenomenon. More specific to the field of optical spectroscopy, a spectrophotometer is the name given to a device that measures the intensity of light as a function of wavelength. In almost all experimental setups, spectrophotometers measure the the intensity of light that is either transmitted through or reflected from a sample.

Although spectrophotometers are complex, fine tuned devices, they all contain three basic components used in the measurement of light: a light source, a diffraction grating, and a light sensor or detector. The light source provides the initial introduction of radiation that is projected onto the sample. It is crucial for the light source to emit a continuous spectrum of radiation that covers the wavelengths

measured by the spectrometer. The diffraction grating takes in a stream of broadband light and refracts the different wavelengths of light at slightly different angles, in essence separating the light so that the sensor in the spectrometer can measure the intensity of individual wavelengths. Lastly, the light sensor or detector is used to measure the intensity of light projected onto it. These three basic components and how they fit into the general construction of a spectrophotometer can be seen in Figure 2.1 below.




Figure 2.1: Basic construction of a reflectance spectrophotometer

The above figure demonstrates the most common form of light sensor on the market which is a linear or two dimensional array of photodetectors positioned to capture and measure light at various wavelengths. The most common style of linear array detectors used in modern spectrometers is a charge-coupled device, or CCD. A CCD is an electronic device composed of an array of photosensors that convert incident photons into electronic charge. The array of sensors similar to pixels on a small screen where each pixel is precisely positioned to measure light intensity at a certain wavelength. The signal from each pixel in the array is individually amplified and used to build up a spectrum of light intensity versus wavelength.

2.1.3 Light Collection Methods

Collecting and projecting adequate light into a spectrophotometer is a necessary physical and mechanical design challenge that is critical in performing quality reflectance measurements. Although there is a wide variety of novel light collection methods used in products across the market, projecting light into a spectrophotometer is generally done in either of two ways. The first method involves projecting the reflected radiation from a source directly into the input slit of the spectrometer. The second involves using fiber optic cables to collect reflected radiation from a source through one end of the cable and then projecting it out the opposite end and into the input slit of the spectrometer. A basic visual depiction of a reflectance spectrophotometer with fiber optic light collection can be seen in Figure 2.2 below.




Figure 2.2: Reflectance spectrophotometer with fiber optic light collection

When utilizing fiber optic cables to collect light, it is often common practice to use a bifurcated fiber bundle, a Y-shaped assembly where two (or more) fiber optic cables are packaged together in a common end and branch out to form two (or more) separate cables at the opposite ends. Bifurcated fibers can be used to route light from a common source to multiple sensors. However, in many spectrophotometers one independent end of a bifurcated cable is connected to a light source while the other is connected to the sensor. Therefore, at the common end, some portion of the fibers will project light onto the sample while the remaining fibers will collect and route reflectance light to the sensor for measurement. A further discussion of fiber optic cables is presented in Section 2.3.

2.1.4 System Requirements

With all of the pertinent foundational information presented above, a basic set of system requirements can now be considered which can be seen in Table 2.1 below. These requirements will drive most of the electrical and mechanical design decisions made throughout the creation of this field reflectance spectrophotometer.

Table 2.1: Basic System Requirements for Field Reflectance Spectrophotometer

Mass	Less than 5 kg
Volume	Less 5,000 cubic centimeters, no more than 30 cm in any dimension
Power	Less than 50 W to operate
Wavelength Range	Must cover at least 400 - 2400 nm
Spectral Resolution	Less than 2 nm in VIS band, 5 nm in NIR band, 10 nm in SWIR band

Cost

Less than \$25,000 for entire system

The maximum mass of the system is mainly determined to allow for practical in field use. Anything above 5 kilograms would result in a system that is too heavy and cumbersome to allow for quick and easy plant measurements. The maximum volume was determined in a similar way, as anything above 5000 cubic centimeters or 30 cm in any dimension would be too bulky to easily transport manually throughout the field. The maximum power is set to limit the draw from the robot's on board battery while also being large enough to power the spectrometers as well as a potentially high power incandescent light source. The wavelength range is set to cover the visible light band (380 - 750 nm), the near infrared band or NIR (750 - 1500 nm), and most of the short wavelength infrared band or SWIR (1500 - 3000 nm). The spectral resolution is set at slightly above the industry standard to allow for high resolution reflectance light analysis. Lastly, a maximum cost is set at \$25,000 to allow for quality components while also forcing the design to remain relatively low cost (compared to some high end sensors) for marketability and economic constraints from this project's funding source.

2.2 Sensor Selection

A overview of spectrophotometers and the necessary requirements for this project were given in the previous section which will form the foundation of the sensor selection process that follows. A wide variety of spectrometers were surveyed, and the primary list that was considered is presented in Table 2.2 below (ordered by price).

Table 2.2: Spectrophotometer Sensors and Specifications

Spectral Sensor	Price	Wavelength Range	Spectral Resolution	Power	Dimensions (mm)	Mass
TI NIRScan Nano	\$1,000	900 - 1700 nm	10 nm	6.5 W	62 x 58 x 32	n/a
Thorlabs CCS200	\$2,868	200 - 1000 nm	2 nm	1.25 W	122 x 79 x 30	400 g
B&W Tek Exemplar	\$2,995	350 - 1050 nm	1.5 nm	2.5 W	102 x 67 x 34	340 g
Ocean Optics FLAME-VIS-NIR	\$3,200	350 - 1000 nm	0.5 nm	1.25 W	89 x 63 x 32	265 g
Ocean Optics FLAME-NIR	\$7,200	950 - 1650 nm	5 nm	1.25 W	89 x 63 x 32	265 g
TI NIRScan	\$8,500	1350 - 2450 nm	6 nm	12 W	112 x 112 x 96	n/a
B&W Tek Sol 1.7	\$10,995	1200 - 1600 nm	1.5 nm	17.5 W	197 x 109 x 68	1400 g
BaySpec InGaAs NIR	\$16,995	1000 - 2000 nm	8 nm	12 W	162 x 105 x 60	650 g
ASD FieldSpec 4 Standard-Res	\$48,000	350 - 2500 nm	1.1 - 1.4 nm	60 W	368 x 292 x 127	5440 g

Finding spectrometers that fit the necessary size and weight requirement proved to be rather difficult as many scientific instrumentation companies offer some form of spectrometer however many of these systems are large, heavy, and designed for laboratory settings. Therefore, the list of products above are all high quality sensors that are also small, portable, and measure light in some portion of the 380 - 2400 nm required wavelength range.

Aside from the price, the ASD FieldSpec 4 Standard-Res is an ideal product that captures the full wavelength range, has an incredibly small spectral resolution, and operates just below the maximum power limit. There is no question that ASD offers state of the art spectral sensors and that their FieldSpec 4 model would be a perfect choice, however the cost is well above the budget, thus a more creative and cheaper solution is needed. Because the ASD spectrometer is the only system that completely covers the necessary wavelength band, multiple spectrometers will need to be selected to operate in unison to provide the full spectral output required.

Visible Light Band

The only remaining sensors that measure light in the visible band include the Thorlabs CCS200, B&W Tek Exemplar, and Ocean Optics FLAME-VIS-NIR. Each of these sensors are very similarly priced, with the highest cost only about \$300 more expensive than the cheapest option. Furthermore, the Ocean Optics FLAME-VIS-NIR has a considerably smaller spectral resolution at three times smaller than the next closest sensor. In addition, the FLAME-VIS-NIR has the smallest dimensions, lowest electrical power, and smallest mass of all three options. It is for these reasons that the FLAME-VIS-NIR was selected for this project.

NIR Light Band

Given that the selected visible light band sensor from above measures light up to 1000 nm, the NIR sensor must cover wavelengths less than or at this value so that there is no gap in the overall reflectance light measurements of the overall system. Therefore, the B&W Tek Sol 1.7 sensor is suboptimal for this application as its wavelength range is 1200 - 1600 nm and also has a relatively large mass and price compared to similar sensors.

The only remaining NIR sensors that meet all requirements are the TI NIRScan Nano, Ocean Optics FLAME-NIR, and BaySpec InGaAs NIR. Although the TI NIRScan Nano has a significantly lower price, the spectral resolution of 10 nm falls well above the required value of 5 nm. Ultimately the Ocean Optics FLAME-NIR was selected over the BaySpec sensor as it provides a smaller specular resolution, is less than half the price, and requires significantly less power to operate. Furthermore, the FLAME-NIR pairs nicely with the previously selected FLAME-VIS-NIR as Ocean Optics provides great software to simultaneously control multiple of their sensors during spectral measurements.

SWIR Light Band

When considering sensors from Table 2.2 that measure light in the SWIR light band, it is clear to see that there is only one sensor remaining, the TI NIRScan. Although there are certainly other sensors on the market that measure light intensity at wavelengths larger than 1500 nm and up to 2400 nm, the vast majority of SWIR spectrometers are designed for permanent use in a laboratory. Simply put, SWIR spectroscopy requires more complex and expensive equipment that usually is not practical for in field measurements. Therefore, the market for high quality, compact, inexpensive, and low power SWIR sensors is small.

The TI NIRScan is ideal for this application as it is a lightweight and low cost sensor that offers a wavelength range of 1350 - 2450 nm with a spectral resolution of 6 nm. The NIRScan also uses a novel technology to collect and record light that makes it especially appropriate and useful for in field applications. Instead of utilizing an array of photosensors such as a CCD, this system contains a digital

micromirror device, or DMD. The DMD is essentially an array of aluminum micromirrors where the angle of each mirror can be controlled to precisely reflect light in certain directions. Various columns in the micromirror array take turns reflecting light at a certain wavelength onto a photosensor while the remaining mirrors are controlled to reflect light away from the sensor. At a very high speed, the columns sequentially reflect light at respectively wavelengths onto the sensor and a full spectrum is built up over a very short period of time. The novel and particularly useful feature of this technology is that the DMD can be programmed quickly and easily to adjust which mirrors correspond to which wavelengths. As a result, the user can control the overall spectral range and resolution of the spectrometer. Due to this inherent flexibility as well as the previously stated benefits of the system, the TI NIRScan was chosen as the best SWIR sensor for this project.

2.3 Light Source

2.3.1 Light Source Background

The selection of a proper light source used in a spectrophotometer is crucial to the accuracy and consistency of spectral reflectance measurements of a sample. At the most basic level, the light source must emit electromagnetic radiation throughout all wavelength bands measured by the sensors in the system. Also of great importance is the intensity, or optical power per unit area, that the light source emits. In many reflectance spectrometers, a significant amount of the radiation originally emitted by the light source is lost due to material absorption, wide angle scattering at the material's surface, and inefficient light coupling through fiber optic assemblies. Regardless of the quality of the system, these forms of light loss are inevitable and thus requires a sufficiently powerful light source to provide reflected radiation with adequate intensity to the photosensors for quality measurements.

Planck's Law and Black Body Radiation

When considering the spectral output and optical intensity of many types of broadband light sources, it is common to model the theoretical performance using Planck's Law and the concept of black-body radiation. A black body is a theoretical body that perfectly absorbs all electromagnetic radiation incident to it. When a black body reaches a constant temperature, and is thus in thermal equilibrium, the body emits radiation based on Planck's Law where the spectral energy distribution is solely determined by the temperature of the body. More specifically, Planck's Law relates the spectral radiance, B_λ (or radiated power per unit angle per unit area per unit wavelength) to temperature as shown here.

$$B_\lambda(\lambda, T) = \frac{2hc^2}{\lambda^5} \frac{1}{e^{hc/\lambda k_B T} - 1} \quad (2.1)$$

Where λ is the wavelength, T is the temperature, h is the Planck constant, c is the speed of light, k_B is the Boltzmann constant. Additionally, Figure 2.3 offers a graphical representation of Planck's Law and the spectral energy distribution of a black body at various temperatures. The theoretical concept of blackbody radiation is best applied to incandescent light sources (a filament that emits light when heated) and the above information is presented to aid in the understand of light source selection throughout this

chapter. There are some common light sources used in spectroscopy applications that are not well approximated by the curves of Planck's Law including LEDs, gas arc lamps, and fluorescent lamps.




Figure 2.3: Black body radiation based on Planck's Law

Color Temperature

Although real world light sources never perfectly emit electromagnetic radiation according to the theoretical bounds of black body radiation, Planck's Law does offer a convenient and relatively similar upper bound to the spectral output of a light source. In fact, the "color temperature" of a light source refers to the temperature of a black body that radiates light most similar to the spectral output of the physical light source. Therefore, the color temperature of a light source has a significant effect on the resulting optical power and wavelength band. A light source with a higher color temperature will be able to deliver higher optical power for less electrical power input while the peak of the spectral power is concentrated in a lower wavelength band, usually visible light. A light source with a lower color temperature produces a spectral output whose peak is centered around a higher wavelength band but usually requires more electrical power to operate due to the fact that black body emitters simply produce less energy at lower temperatures. For these reasons, it is common to see light sources with very large color temperatures used with visible light spectrophotometers while light sources with much lower color temperatures are needed for near infrared and infrared spectrometers.

Common Bulb Types in Spectroscopy

LED - A unique type of light source consisting of a diode that emits light due to the electroluminescence effect. The diode consists of two adjacent semiconducting materials, one with an excess of electrons and one with a lack of electrons. When current flows across the diode, electrons are able to go from the conduction band of the negative material to fill the holes in the valence band of the positively charged semiconducting material which releases energy in the form of photons. Because the frequency of emitted photons is determined by the discrete energy difference between these two electron states, the emission spectrum of an LED is usually very narrow, almost monochromatic, and concentrated somewhere in the visible light band. However, it is common to achieve more broadband spectral emission from LEDs through the use of phosphor coatings. Phosphors are materials that absorb radiation at a

certain wavelength and then emit light at another wavelength through a process called fluorescence. For example, a common way to produce white light from an LED is with a InGaN blue LED and a phosphor coating that absorbs some of the blue light and emits yellow light. The yellow light from the coating and the original blue light combine to create a white light that covers much of the visible light band.

Deuterium - A common bulb used in spectroscopy applications where the measurement of ultraviolet light is critical. This type of light source is a gas discharge arc lamp that emits radiation from about 100 nm to 900 nm with the main, continuous spectrum of radiation concentrated in the ultraviolet band at 250 to 300 nm. The bulb is filled with deuterium gas, a stable isotope of hydrogen, at a low pressure that emits UV radiation when the molecules are excited to a higher energy state by an electric arc is produced between two electrodes inside the bulb. Deuterium arc lamps are preferred over similar hydrogen lamps as they produce a wider spectral output in the UV band and also have a much longer lifetime. However, Deuterium arc lamps are much more expensive than most other bulb types.

Xenon - A type of high-intensity discharge lamp with great emission in the ultraviolet band. Similar to deuterium arc lamps, these light sources consist of a glass tube filled with ionized xenon gas at a high pressure that emits electromagnetic radiation when an arc of electricity is passed between two electrodes at each end of the bulb. Most xenon lamps operate at a color temperature around 5500 to 6000 K resulting in a bright white light and a spectral output where the majority of light intensity is focused between 250 nm to 1000 nm but also continues up to 2500 nm at significantly lower intensity. Xenon arc lamps are common in spectroscopy given their high UV emission, however they suffer from the fact that they require a large amount of input electrical power.

Tungsten - A common type of incandescent bulb where light is emitted from a tungsten filament by heating it to very high temperatures through electric current. Because these light sources rely on the principle of incandescence, they are very inefficient (more than 90% of input energy is converted to heat) and require much more power compared to other sources to produce the same amount of light. However, tungsten filament bulbs are well approximated by the principles of black body radiation and thus tend to produce a very continuous, wide spectrum across the visible and infrared wavelength bands. This fact makes these sources a common selection in infrared spectroscopy applications. Additionally, the filament is almost always encased in a glass or quartz bulb to prevent oxidation of the tungsten. Furthermore, it is common to fill the bulbs with low pressure halogen gas. The gas creates a phenomenon known as the Halogen cycle where the gas particles will deflect evaporated tungsten atoms back onto the filament to improve its lifespan and also allow operation at higher temperatures. These sources are often called halogen lamps or tungsten halogen lamps.

2.3.2 Light Source Selection

There are many light source assemblies on the market that are designed and optimized for spectroscopy applications. These assemblies come equipped with a bulb or lamp of some kind, a series of optical lenses to direct and focus the light appropriately, and hardware to manage the input electrical power and output optical power. The main advantages of using an off the shelf light source include steady optical output, robustness, reliability, and no need for mechanical or electrical design work. However, these assemblies are often expensive and rather large when considering the strict size constraints on this spectrometer project. Thus it was ultimately decided to design a custom light source for the reflectance spectrometer system.

To give a sense of the variety of light sources on the market, Table 2.3 below presents a list of the light source assemblies considered for this custom reflectance spectrometer application.

Table 2.3: Light Sources and Specifications

Light Source	Bulb Type	Price	Wavelength Range	Bulb Electrical Power	Lifetime	Dimensions (mm)
Thorlabs SLS201L	Tungsten Halogen	\$1,030.20	360 - 2600 nm	9 W	10000 hrs	216 x 57 x 55
Thorlabs SLS202L	Tungsten	\$1,199.52	450 - 5500 nm	7.2 W	10000 hrs	216 x 57 x 55
Thorlabs SLS203L	Silicon Nitride Globar	\$1,596.30	500 - 9000 nm	24 W	10000 hrs	209 x 57 x 55
Ocean Optics HL-2000-LL	Tungsten Halogen	\$850.00	360 - 2400 nm	4.75 W	10000 hrs	150 x 62 x 60
Ocean Optics HL-2000-HP	Tungsten Halogen	\$950.00	360 - 2400 nm	20 W	1000 hrs	150 x 62 x 60
Newport 66088	Quartz Tungsten Halogen	\$1,207.00	350 - 2100 nm	150 W	200 hrs	241 x 184 x 119
Edinburgh Instruments SL1	Tungsten Halogen	\$375.00	350 - 2200 nm	4 W	10000 hrs	90 x 76 x 38
B&W Tek BPS101	Tungsten	\$795.00	350 - 2500 nm	5 W	5000 hrs	102 x 63 x 41
B&W Tek BPS2.0	Tungsten	\$1,495.00	350 - 2600 nm	20 W	2000 hrs	121 x 80 x 67
B&W Tek BDS130A	Deuterium and Tungsten	\$2,795.00	190-2500 nm	70 W	2000 hrs	155 x 76 x 64

The average cost among all of the light sources listed is \$1,229.30 which is indeed a small percentage of the overall budget. However, considering the total cost of the three spectrometers at \$18,900 and the remaining costs of the fiber optics and electrical and mechanical components, purchasing an off the shelf light source would consume a large portion of the remaining budget and remove flexibility in future design decisions. Furthermore, the average volume of the above light sources is 1032.8 cubic centimeters which consumes over 20% of the overall volume constraint. The volume of the three spectrometers is about 1563 cubic centimeters or about 31% of the volume. Additionally, significant space must be allocated for components such as USB hubs, cooling devices, connectors, and general vacancy needed to prevent overcrowding and overheating. Although an off the shelf light source could most likely be cleverly packed into the overall spectrometer housing, it again limits the flexibility and creativity for future design decisions. For these reasons, it was determined to create a custom light source assembly to achieve similar performance at a lower cost and smaller volume.

Custom Lamp Selection

The ensuing selection of a bulb/lamp was mainly driven by performance and reliability. Knowing that this reflectance spectrometer needed to operate at wavelengths from 400 nm up to 2400 nm, a tungsten bulb was the clear choice, given its continuous spectral output across the visible and infrared bands. The main challenge in creating a custom lighting source was ensuring that the filament of the lamp

was near perfectly aligned with the optical lens that would direct and focus the light onto the end of a fiber optic cable. Given the incredibly small area of a fiber optic cable, 0.79 mm^2 for a $1000 \text{ }\mu\text{m}$ diameter fiber, any slight misalignment on the order of a fraction of a millimeter would significantly reduce the light throughput through the fiber and onto the specimen. It is for this reason, that selecting a lamp that comes precisely mounted by the manufacturer was of great importance.

Ultimately, the Thorlabs HEP3965 Broadband IR Tungsten Bulb was selected, shown in Figure 2.4 below. This source produces a broadband output from about 400 to 5000 nm at a color temperature of 1900 K, resulting in a peak intensity right around 1500 nm. The bulb comes mounted in a small TO-8 can with a reflector placed behind the filament for greater forward facing luminous flux. The bulb can be driven up to 3.5 V with a maximum current of 2 A resulting in a maximum power rating of 7 W.

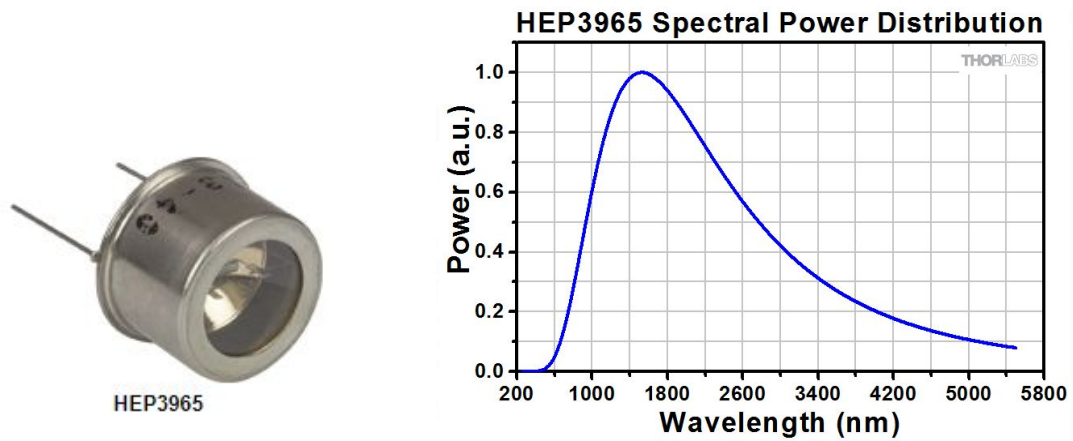


Figure 2.4: Thorlabs broadband IR tungsten bulb and spectral output [29]

As will be discussed further below, this bulb comes with a threaded mounting adapter allowing for easy installation into lens tubes. These components allow for greater precision in filament alignment and imaging onto the end of the optical fiber.

2.4 System Overview

2.4.1 Custom Light Source Design

In addition to the previously selected bulb, the overall lighting assembly must contain optical lenses to focus the light onto an attached fiber optic cable, a connector for the fiber optic cable, and a rigid housing to contain all of these components. The overall assembly can be seen below in Figure 2.5 along with an exploded view and component overview in Figure 2.6.



Figure 2.5: Custom light source assembly




Figure 2.6: Light source exploded view

Detailed Parts List:

- (a) **Thorlabs S120-SMA** - Lens tube cap with internal SM1 threading and SMA fiber connector
- (b) **Thorlabs SM1RC and TR1** - SM1 Lens tube slip ring and vertical mounting post
- (c) **Thorlabs SM1S20** - 2" Long SM1 lens tube spacer
- (d) **Thorlabs CMRR** - Retaining ring for SM1 lens tube
- (e) **Thorlabs LB1761** - N-BK7 Bi-Convex Lens with 1" diameter and 25.4 mm focal length
- (f) **Thorlabs LB1761-ML Housing** - Mounting tube for lens and retaining ring
- (g) **Thorlabs HSLT2** - Passive heat sink lens tube with internal SM1 threading
- (h) **Thorlabs HEP3965** - Broadband infrared tungsten bulb
- (i) **Thorlabs HEPM** - Mounting adapter for HEP3965 bulb

As seen in the figures above, the tungsten bulb (h) rests inside a thick, circular adapter plate (i) with external threading. This adapter plate is then mounted inside a heat sink style lens tube (g) with many radial fins for added heat expulsion and greater convective cooling during use. Because the interior

of the heat sink lens tube and the exterior of the bulb adapter plate are both finely threaded, the position of the bulbs filament along the axis of the tube can be precisely controlled. Attached to the front face of the heat sink lens tube is a short mounting tube (f) for an N-BK7, a type of borosilicate glass, biconvex lens (e). The next component along the optical axis is a two inch long lens tube (c) to properly space the lens from the optical fiber. On the end of the lens tube spacer is a threaded cap (a) with an SMA fiber connector to attach a fiber optic cable for light collection. The entire assembly is held in space by a vertical mounting post with a large slip ring (b) that tightly hugs the two inch lens tube. This vertical mounting post can be fastened to the base plate of the spectrometer housing for rigid attachment within the larger system.

The total cost of this custom light source assembly is \$375.66 which is about 30% of the average cost of the off the shelf products previously presented. Furthermore, the total volume (of a bounding box around the assembly) is 335.16 cubic centimeters with dimensions of 126 mm in length, 75.5mm in height, and 56 mm in width. The overall volume is 32% smaller than the average volume of the off the shelf products.

2.4.2 Mechanical and Electrical System Design

Mechanical Design

The overall reflectance spectrometer assembly is responsible for housing the three spectrometers, the custom light source, and all relevant electronics such as a USB hub, power regulators, and cooling systems. The aluminum housing is 26 cm in length, 16.5 cm in width, and 11.5 cm in height for a total volume of 4933.5 cubic centimeters, or just below the maximum volume specified. The overall mass of the system is about 3.5 kg which is well below the maximum mass specified. A final rendering is shown below in Figure 2.7 and an interior view is shown in Figure 2.8.



Figure 2.7: Reflectance spectrometer system




Figure 2.8: Reflectance spectrometer interior:

(a) Custom broadband light source; (b) Ocean Optics FLAME-VIS-NIR spectrometer; (c) Ocean Optics FLAME-NIR spectrometer; (d) Ethernet and USB bulkhead connectors; (e) TI NIRScan spectrometer

The overall system housing is built out of six aluminum panels, which account for about 33% of the total mass, bolted together with M3 x 10 socket head cap screws that are slightly recessed into the faces of each panel to create a smooth, flat surface on all sides. Given the need for easy access to the fiber optic inputs of the three spectrometers from the exterior, two separate 3D printed adapter plates are mounted to the front panel of the system that protrude into the interior of the housing with small openings for the three SMA fiber optic connectors. The fiber optic connector for the light source is accessible through a small hole on the left side panel of the housing. USB and ethernet connections are provided through two rugged and sealable bulkhead connectors on the front panel. The TI NIRScan spectrometer is digitally interfaced over ethernet while the two Ocean Optics spectrometers are accessed over USB.

A sealed four pin electrical connector is mounted to the back panel of the housing to provide power to the entire system. Additionally, a rocker switch is positioned next to this connector to turn on and shut off power to the spectrometer. Two fans are also mounted to the inside face of the back panel to provide convective air flow to cool the heat sink around the light source. Similarly, two fine filter screens are placed on the outside face over the fans to prevent debris from entering.

Electrical Design

Two DC-DC converters are mounted underneath the TI NIRScan to provide regulated power to the subsystems within the housing. A CUI converter (PYB20-Q24-S3-T) with a 9-36 V input range and 3.3 V, 5 A output is used to power the custom light source. The output voltage is trimmed up to 3.5 V using a resistor in line between the trim and Vo+ pins on the converter. The tungsten bulb is the only load powered off of this supply as to mitigate noise and disturbances in the output signal. A circular push

button switch is located on the left panel above the SMA connector to the light source to control power to this DC-DC converter. Additionally, a second CUI converter (PQA30-D24-S12-T) with an 18-36 V input range and 12 V, 2.5 A output is used to power the remaining subsystems; the TI NIRScan spectrometer, a USB hub to power the two Ocean Optics spectrometers, the two cooling fans, and a small LED on the light source switch.

The four pin power connector on the back panel splits to provide positive and ground lines to each of the two converters. The set of wires feeding to the 12 V output converter is designed to receive ~24 V unregulated power straight from the battery pack on the robotic platform. The set of wires feeding to the 3.5 V output converter has an input range of 9-36 V so it can take in either ~24 V unregulated power or power from the 12 V regulator on board the robot.

2.4.3 Software

This reflectance spectrometer system is largely designed to work directly with the agricultural robotic system described in Chapter 1. Given that this robotic system uses the Robotic Operating System (ROS) as the main software framework, custom software was written in C++, and wrapped into ROS nodes, to provide low level communication with the three spectrometers.

The TI NIRScan has an on-board processor that runs an open source spectrometer application on a Linux kernel. Usually, this application is accessed by the user through a web page and controlled through an RNDIS connection over USB or Ethernet on the local area network. Therefore, automated data collection using the TI NIRScan is controlled using the libssh library to set up a simple, automated client server communication using the SSH protocol and the IP address of the device. The Ocean Optics spectrometers are normally controlled through a desktop application called OceanView. However, the company also offers an open source C/C++ device driver library called SeaBreeze which was used to create software to interface with these devices and automate the data collection process using custom ROS nodes.

Chapter 3

Design of the Leaf Grasping Manipulator

3.1 The Existing Stalk Grasping Manipulator

System Overview

The first robotic manipulator on the mobile robotic platform was designed to detect and grasp sorghum stalks. The end effector consisted of two aluminum fingers that when actuated would enclose the stalk against the end of the arm. A rind penetrometer, or small needle connected to a force gauge, was then applied to the sorghum stalks. The force to penetrate the stalk wall with this needle was then recorded and used to predict stalk strength, a key trait used by plant geneticists in improving overall yield of the crop. The entire robotic arm and custom stalk grasping manipulator was designed by Merritt Jenkins, a past graduate student in robotics at Carnegie Mellon, whose works is detailed in his Master's thesis [26]. Jenkins provides a thorough description of the mechanical design, controls, and components required to build his custom manipulator in Chapter 4 of his thesis; the key points are summarized below. Renderings of the forearm link and end effector of the stalk grasping manipulator are also shown in Figure 3.1 below.




Figure 3.1: Stalk grasping manipulator

System Requirements

The key system requirements considered when designing this custom manipulator were as follows. The arm must be able to grasp plants up to 0.5 meters away. The total system must not exceed

100 W of power to operate. The overall mass of the system must be less than 25 kg while being fit for outdoor operation. Furthermore, the arm must have the capacity to move the end effector (up to 2 kg in mass) at a speed of 1 m/s. Lastly, the entire system must cost less than \$20,000. Due to this strict set of requirements and the unique outdoor, agricultural setting, a custom manipulator was built to keep costs low while also providing absolute control in system design. The degrees of freedom needed from any manipulator is largely based on the complexity of motion required and ultimately has a significant effect on the overall mass and form factor of the assembly. For this challenge of grasping sorghum stalks and applying a rind penetrometer, it was determined that the end effector could be applied to any point along the circumference of the stalk. Therefore, only three degrees of freedom were required; one prismatic joint to position the gripper vertically in space and two revolute joints to position the gripper to any location on the horizontal plane agnostic of angle.

Hardware and Sensor Selection

Three HEBI X5 actuators were selected as the motor modules to control all three degrees of freedom. Two of these actuators are mounted at the shoulder and elbow joints of the robotic arm and act as revolute joints to horizontally position the end effector. The third motor actuates the driveshaft of a Thomson Linear Motion Movopart M75 linear stage mounted vertically on the top plate of the mobile robotic platform. These motors were selected due to their small volume and slim form factor, impressive IP rating for outdoor performance, and series elastic actuation. The end effector makes use of a Maxon EC16 BLDC motor with a Maxon GP16S ball screw to produce the linear motion required to drive the rind penetrometer needle and actuate the two fingers of the gripper.

In terms of visual sensor selection, a DUO3D MLX stereo camera was initially mounted to the end of the forearm to perform stalk detection and visual servoing. As the importance of visual servoing decreased and the need for higher resolution, dense point clouds increased, a Carnegie Robotics MultiSense S7 was purchased and mounted rigidly to the vertical mast of the robot.

3.2 Adaptation for Stalk Spectroscopy

Given that a custom manipulator for grasping sorghum stalks was already developed by Jenkins with a proven viability in the field, the next logical step was to add the ability to apply the reflectance spectrometer to sorghum stalks in addition to the rind penetrometer. Therefore, a simple update to the mechanical design of the end effector was made to mount a bifurcated fiber optic cable to the manipulator. The updated version of the original stalk grasping manipulator can be seen below in Figure 3.2.




Figure 3.2: Updated stalk grasping manipulator with fiber optic attachment

The red 3D printed V-block on the inside of the gripper, originally seen in Figure 3.1, was adjusted to include a small vertical wall through which a telescoping set of spring loaded cylinders is mounted. The end of the fiber optic cable is threaded into a small SMA bulkhead connector mounted inside the thinner, white cylinder. The small grey cap on the end of this cylinder protrudes slightly past the inside of the V-block but is spring loaded so that it collapses inward when a stalk is grasped to ensure firm contact. As previously stated, this configuration makes use of a bifurcated cable where the single cable is split into two fiber legs, one from the light source and the other from the TI NIRScan. At this stage in the project, the two lower wavelength Ocean Optics spectrometers were not a part of the reflectance spectrometer system. Initial testing occurred in the sorghum test plots in Pendleton, South Carolina in the summer of 2017, the details of which are fully discussed later. Even in these initial tests, this updated manipulator proved to provide spectral scans of sorghum stalks with the same reliability as manual scans while improving the efficiency of data collection.

3.3 Robotic Leaf Grasping and Requirements

3.3.1 Problem Statement

In field crop analysis is certainly not constrained to the stalks. In fact, a wide range of research in plant genetics, and agriculture in general, relies on applying contact sensors to the leaves of plants. Frequent contact based sensors include small mechanical devices that measure leaf thickness, leaf wetness sensors used to measure moisture levels and water deficit stress, and portable spectrometers used to measure spectral reflectance of leaves for chemical and compositional predictions. Given that a portable, infield spectrometer has been designed for this project, the logical next step is to use a robotic manipulator to gather reflectance measurements from leaves. Therefore, the goal of this section is to develop such a manipulator to autonomously grasp sorghum leaves and gather spectral reflectance data. The following discussion details the design and development of a custom leaf grasping mechanism to accomplish this task.

3.3.2 System Requirements

When the notion of robotic leaf grasping was first proposed within the scope of this project, it was immediately made clear that any new manipulator must be heavily derived from the previous stalk gripping manipulator. This was mainly the case to keep costs low while also limiting any significant change to the design of the mobile robot base. A functioning custom robotic arm along with software to control it through ROS was already in place, so recycling as much of this work into a new leaf grasping manipulator was seen as the best solution. Furthermore, if executed properly, the two manipulators could be quickly and easily swapped in and out depending on the task. With this challenge in mind, a list of system requirements were created to guide the design and development of this custom manipulator which can be seen in Table 3.1 below.

Table 3.1: Basic System Requirements for Leaf Grasping Manipulator

System Mass	Less than 20 kg
End Effector Mass	Less than 2 kg
Reach	Must grasp leaves 0.5 m away
Power	Less than 100 W to operate
Sealing	Must be rated for continuous outdoor use
Cost	Less than \$1,000 for updates to stalk gripper

The overall system mass was determined mainly by the original constraint set in the design of the stalk grasping manipulator. Extensive calculations were made by Jenkins [26] to find a system mass and geometry that would ensure safe operation of the vertical stage motor and revolute joint actuators even under worst case dynamic loading during field testing. This also further enforces the constraint that any alterations or additions to the original robotic arm design must be kept to a minimum. Similarly, the new end effector must be under 2 kg in mass.

The reach capability of 0.5 m easily encompasses leaves to the immediate sides of the robot given the ~ 0.38 m distance between the centerline of the robot and the sorghum rows. More importantly, this reach allows for the ability to grasp leaves in the space behind the robot up to and just past the back edge of the robot base. Similar to the original manipulator power constraint, the new design must take less than 100 W to operate in order to maintain current battery life and to not overload the max power ratings of the onboard DC-DC converters. The entire arm assembly and end effector must be properly sealed to survive continuous and long term outdoor use. As such, the manipulator should survive exposure to dust and dirt as well as light splashes. Lastly and potentially most challenging, the updated leaf grasping design must accomplish its goal under all the above requirements and constraints for less than \$1,000.

3.4 Custom Manipulator Overview

This section provides an exploration into the mechanical, electrical, and software design of a custom manipulator for automated leaf grasping. As previously stated, a large part of this manipulator is

heavily derived from the stalk grasping manipulator that was first developed for this project, and as such, the main novelty here is the specialized end effector design. This custom manipulator can be seen in Figure 3.3 below.



Figure 3.3: Leaf grasping manipulator

3.4.1 End Effector Design

Given that the primary goal driving the design of this custom manipulator is to grasp leaves, the physical grasping mechanism was the first aspect of the system to be considered. The most common agricultural manipulators usually make use of pneumatic devices to secure an object through suction, mechanical fingers that enclose and secure the target object, or some combination of these two techniques. Given that end effectors with suction cups are common in many automated fruit harvesting applications, using some form of suction to perform leaf grasping was first investigated. A rough prototype consisting of a long, flexible tipped suction cup actuated by a servo motor was mounted to the end of the pre-existing robotic arm. To test this design, the robotic arm was placed in an environment surrounded by natural and synthetic leaves. After many failed attempts to grasp and secure the leaves, it became clear that using suction cups for this application would not work well. Even with the long, flexible tip of the suction cup, the target object needs to have enough mass and thus inertia to provide an adequate resistive normal force so that the rubber tip of the cup can deform enough to create a seal around the object. When the end effector was pressed up against a leaf, the leaf would easily deform away from the suction cup resulting in an inability to create suction between the two objects. These tests made it clear, however, that some sort of mechanical device was needed to physically surround the leaf for successful grasping.

Therefore, the design for a new end effector was created that consists of two wide pads that clamp together around a leaf. More specifically, the gripper resembles a sort of jaw mechanism where the upper plate is actuated by a servo motor and the bottom plate is rigidly attached to the frame of the arm. When the end effector is positioned around a leaf, the top plate rotates downward to firmly press the leaf against the rigid bottom plate. A closer view of the end effector can be seen in Figure 3.4 below.



Figure 3.4: Custom leaf clamp end effector

This manipulator is also responsible for collecting reflectance spectra of leaves upon grasping. Therefore, another key component of the end effector is the way in which it positions fiber optic cables directly above the surface of leaves for collection of reflected light. From Figure 3.5 below, it can be seen that the upper red pad of the clamping mechanism has a rectangular opening near the top for mounting a separate, small adapter plate used to accurately position up to four fiber optic cables. The four cables are equally spaced around the circular opening on the flat, bottom face of the adapter plate so that the axis of each cable intersects at a common point in the center of this circular opening. Additionally, each cable sits at an angle of 45 degrees off of the surface normal of a grasped leaf. Lastly, the exposed end face of each fiber is positioned at an axial distance of 2 mm away from the center of the circular opening.




Figure 3.5: Fiber optic alignment on the end effector

In order to provide feedback to the system on whether a leaf was successfully grasped and a spectral scan should be performed, an LED is mounted in a small cylindrical housing in the middle of the upper red clamp plate which can be seen in the right image of Figure 3.5. A light sensor is mounted in a small opening along the surface of the rigid bottom plate so that the LED and sensor are positioned directly opposite of each other when the plates of the leaf clamp are touching. Therefore, if a leaf is between the two plates, the light sensor will detect little to no intensity from the LED and a successful grasp can be determined.

3.4.2 Mechanical Design

Just as with the stalk grasping manipulator, this system retains the same three degrees of freedom: a vertically mounted linear stage acts as a prismatic joint perpendicular to the ground plane and two HEBI X-5 actuators function as revolute joints at the shoulder and elbow of the two link robotic arm. In order to satisfy the reach requirement of 0.5 m, the aluminum tubes that make up the two arm frames were designed to provide an overall distance of 0.66 m from the axis of rotation of the shoulder joint to the center of the clamping plates on the end effector. Furthermore, the entire leaf grasping manipulator has an overall system mass of about 2.7 kg which is significantly less than the 20 kg constraints placed on the design. This low mass is largely due to the use of thin walled aluminum tubing to form the frame of the robotic arm as well as the use of 3D printed ABS plastic components to make up the end effector. In fact the entire end effector design has a mass of just under 0.4 kg which again is significantly less than the 2 kg constraint detailed above. The mechanical components used to actuate the upper plate of the clamping mechanism can be seen in Figure 3.6 below.




Figure 3.6: Leaf clamp actuator design:

(a) Tower Pro MG995R servo motor; (b) Servo motor shaft coupler; (c) Flange mount bearings; (d) 3/16” steel shaft; (e) Upper clamp plate; (f) Shaft mounted clamp plate support; (g) Nylon bevel gears; (h) Servo shaft support

A Tower Pro MG995R high torque servo motor was selected to actuate the upper clamp plate. Although servo motors do not typically provide an incredibly robust and durable system for outdoor use, they do provide a low cost, light weight solution for actuating small loads with position feedback. For this application, significant effort was placed on creating a robust seal around these electromechanical components, thus the use of a servo motor seemed appropriate. This servo provides up to 8.5 kg-cm of torque at 5 V with a range of about 170 degrees. A shaft coupler connects the typical star toothed output shaft of the servo to a 3/16” diameter steel shaft that is secured by a shaft support to prevent any off axis bending forces on the servo motor. A set of small, nylon bevel gears transfers the rotary motion 90

degrees to another perpendicularly mounted 3/16" steel shaft. This shaft is mounted on each end by flanged ball bearings that are secured to the end of the aluminum frame of the forearm. Two custom shaft mounts were designed and machined to fit the geometry of the system and mate the upper clamp plate to the shaft.

In order to provide an adequate, semi-weatherproof sealing as specified by the system requirements, the open ends of each aluminum tube frame were closed off with some sort of plastic block. The frame between the shoulder and elbow joints is sealed on each end using delrin block inserts as best seen in Figure 3.3. The aluminum frame that holds the end effector is sealed using a similar delrin block insert on the end opposite the leaf clamp. In order to easily access the electronics and mechanical components within this section of the arm, a large cutout was made on the top face of the aluminum tube. This large opening is sealed using an ABS plastic plate that fully covers the opening and extends down the sides of the frame. Lastly, the opening in the aluminum frame at the leaf clamp end was sealed using a custom fit, 3D printed cap insert that is best seen in Figure 3.4 above. This cap is also designed into the rigid base plate of the leaf clamp in order take advantage of the ease of complexity that comes with 3D printing as well as to reduce the number of overall parts and fasteners in the assembly.

3.4.3 Electrical Design

The entire manipulator system can be powered off of 18-36 V and consumes an average of 45 W of power. The HEBI X-5 actuators draw an average current of 0.8 A at 24 V with a peak current draw of 2.4 A. On board regulation allows for these motors to be powered by 18 - 48 V. The servo motor is powered off of 5 V, which is provided by a CUI PDQ30-Q24-S5-D DC-DC converter that steps down a 9-36 V supply to 5 V for the servo. Under maximum loading, the servo draws up to 1.2 A of current. All together, each subsystem in the manipulator can be powered off of the unregulated ~24 V supply from the robot's battery.

An Arduino Nano mounted on a small board behind the servo motor is used to control the motor as well as the state of the LED and analog voltage readings from the light sensor. A long USB cable is fed through the two aluminum frames and down to a bulkhead connector for communication with one of the onboard NUC computers. In a similar fashion, the HEBI actuators communicate with the onboard computers through an Ethernet cable that is routed through the robotic arm and down to the robot base.

3.4.4 Software

Similar to the software design of the spectrometer, the leaf grasping manipulator must be controlled through ROS. HEBI, the company that manufactures the shoulder and elbow actuators, provides great ROS support and a very nice API for custom software development. Therefore, ROS nodes that publish and subscribe to joint state topics control the motion of the three degrees of freedom of the robotic arm. Motion control is largely done through a PID controller with gains set and tuned through a GUI application called Scope also provided by HEBI. The Arduino is interfaced and controlled with ROS through the use of the `rosserial_client` package. The Nano runs software that creates a simple subscriber to a topic that contains a joint angle for the leaf clamp and uses a PWM signal with angular feedback to position the upper clamp plate at the desired angle.

Chapter 4

Leaf Detection and Manipulation

4.1 Leaf Detection

4.1.1 Pipeline Overview

This section details the computer vision pipeline that begins with digital images of sorghum captured on board the robot and quickly and reliably outputs a simplified, three dimensional model of the stalks and leaves within the scene. As stated above in Chapter 3, the Multisense S7 stereo camera by Carnegie Robotics is rigidly mounted to the side of the vertical mast of the mobile robot base and provides dense 3D range data of the environment. The Multisense S7 begins by capturing simultaneous images from right and left cameras and uses an onboard FPGA to perform rapid correspondence matching between images to calculate distance values for points in the scene. Through this process, the Multisense can output depth maps and colored point clouds at up to 15 fps. An example of a set of stereo images from the Multisense as well as the resulting point cloud can be seen below in Figure 4.1.

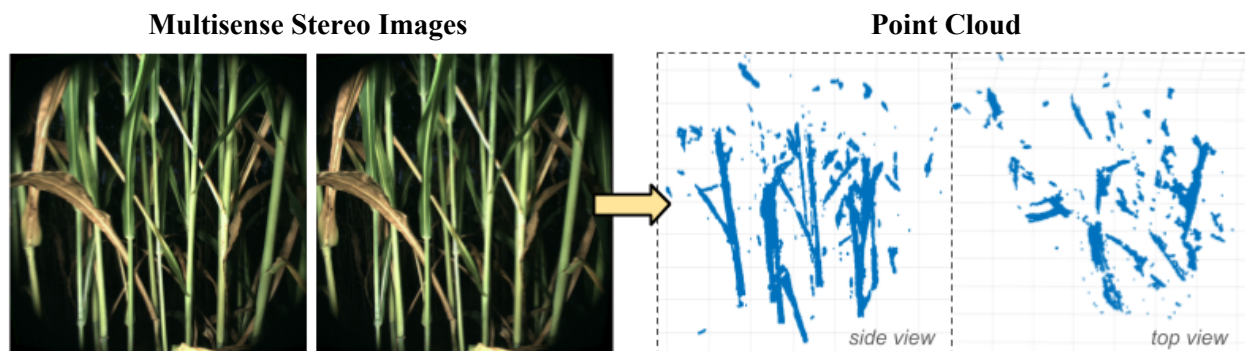


Figure 4.1: Example of stereo images and resulting point cloud

Since three dimensional visualization of the plants is made rather easy with this stereo imager, the pipeline essentially begins with processing the raw point cloud. It should be noted that many of the images within the figures of this chapter are created using MATLAB. This is due to the fact that initial algorithm design and testing was done using MATLAB to simplify complex mathematical operations and also quickly and easily visualize steps within the pipeline. However, the core algorithms were eventually recreated in C++ with the use of the Point Cloud Library (PCL) for speed and functionality through ROS.

PCL is a large, open source collection of algorithms used for image and point cloud processing. A full visual depiction of the leaf detection pipeline is shown below in Figure 4.2.

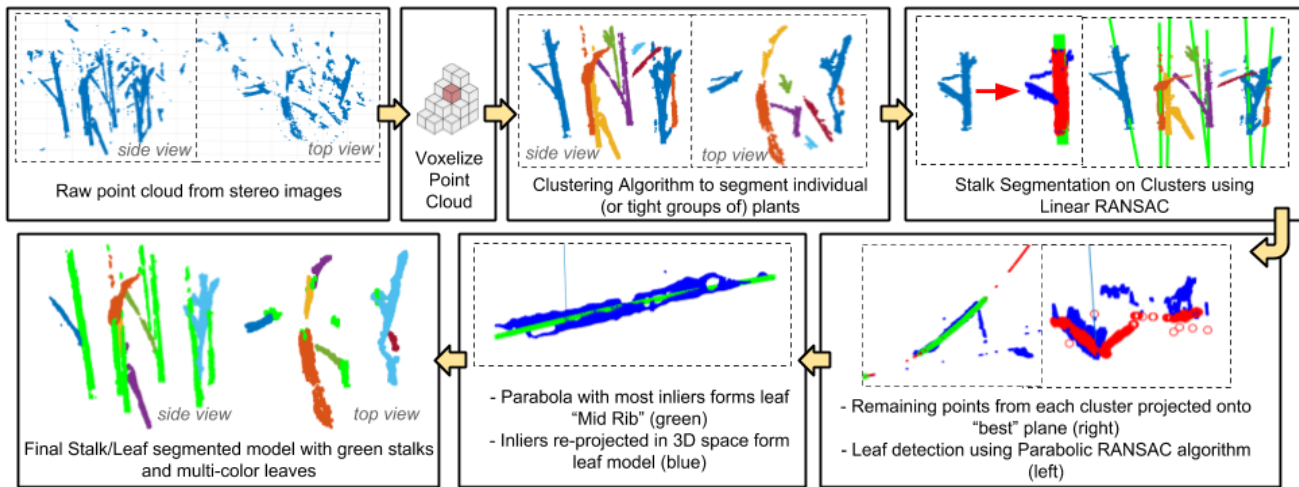


Figure 4.2: Leaf Detection Pipeline

The leaf detection pipeline begins by creating a voxelized representation of the raw point cloud. This method significantly reduces the number of points in the cloud by instead storing only the centroids of points within small discretized 3D volume elements. This voxel grid is then converted into a k-d tree using PCL, a form of binary search tree that dramatically speeds up searching for points throughout space. A clustering algorithm is then employed, using the k-d tree for added efficiency, to detect and store large groups of points using a 2 cm distance threshold. Only clusters consisting of a certain number of points, in this case 100 points, are saved for later processing.

The next stage of the pipeline works to identify stalks within these large clusters through the use of an iterative linear random sample consensus (RANSAC) algorithm. Ultimately, all points belonging to stalks will be segmented out of the cloud in order to simplify the search for leaves. The linear ransac algorithm attempts to find one or more 3D line(s) that best fit the data within each cluster while ignoring potential outliers. This linear RANSAC algorithm for stalk segmentation is discussed in much more detail in section 4.1.2 below. For each cluster, all of the points that are not determined as inliers to a detected stalk are retained while the stalk inlier points are segmented out.

At this point in the pipeline, the clusters now represent a significantly downsampled version of the original point cloud where all stalk points as well as noisy outlier points have been removed. Thus, the existence of leaves should now be much easier to detect. Next, another RANSAC style algorithm is used that attempts to fit parabolic leaf models to the remaining points. A second order polynomial functions as an accurate model for sorghum leaves given the natural curvature. The details of this parabolic RANSAC leaf detection algorithm are discussed in much more detail in Section 4.1.3 below. The inliers of a successful parabolic leaf model are considered to be points that belong to a leaf in the scene. Furthermore, the inliers projected onto the parabolic model form a clean parabola in 3D space that cuts through the middle of each leaf to form a sort of virtual leaf midrib. The inliers from both the linear stalk detection and parabolic leaf detection algorithms can be plotted together to form a clean 3D model

of all the sorghum stalks and leaves from the original scene, as shown in the bottom left image in Figure 4.2.

4.1.2 Basic Stalk Segmentation with Linear RANSAC

As stated in the previous section, stalk segmentation is performed using a RANSAC algorithm that attempts to fit a line to each cluster by maximizing the number of points within a certain threshold distance while ignoring outliers in the data. The random nature of a RANSAC algorithm makes it non-deterministic, however, it is often run many thousands of times to ensure an accurate model is found. The steps of the linear RANSAC stalk detection algorithm are detailed below and the key parameters used throughout the algorithm can be seen below in Table 4.1.

Linear RANSAC Stalk Detection Algorithm

1. Two points within the cluster are randomly selected to uniquely define a line in 3D space and the model parameters ($ax + by + c = 0$) are derived and stored.
2. The angle between the the normal vector of the ground plane and the line is found. If the angle is greater than the *vertical angle threshold* then the current iteration is over and step 1 is repeated.
3. For all other points in the cluster, find the Euclidean distance to the line. If the distance is less than the *distance threshold*, then the point is stored in a current list of inliers.
4. If the number of inliers is greater than the number of inliers from the previous best model as well as the specified *minimum model size*, then the current model is stored as the best model.
5. Repeat steps 1-4 until the *total number of iterations* has been reached.
6. The model parameters of the line with the most inliers are used to represent a stalk. The algorithm is repeatedly run on the set of outliers from the previous result until there are no more stalks found in the cluster that meet the *minimum model size* constraint.

Table 4.1: Key Parameters in the Stalk Detection Algorithm

Number of Iterations	10,000
Vertical Angle Threshold	10 degrees
Distance Threshold	1 cm
Minimum Model Size	75 points

4.1.3 Detecting Leaves with Parabolic RANSAC

The Parabolic RANSAC leaf detection algorithm is similar in principle to the linear RANSAC method described above. In essence, the algorithm attempts to find the set of model parameters, in this case for a parabola, that maximize the number of inliers in a cluster of points. However, the key difference with this algorithm is in the complexity of solving for the model parameters given a random set of points. Intuitively, it may seem as though three randomly selected points would define a unique parabola in three dimensional space much like three points can define a parabola in two dimensional space, however, this is far from the case. In two dimensions, three points of the form (x,y) only uniquely define a parabola given the constraint that the parabola is a function of y in terms of x, and thus has an

axis of symmetry parallel with the y-axis. Without this constraint, a set of three points can be fit by an infinite number of parabolas in most cases. Therefore, in order for a set of three points to uniquely represent a single parabola, the plane on which the parabola will lie as well as a vector specifying the direction of the axis of symmetry must be given.

The general shape and orientation of sorghum leaves can be used to help form these constraints. Primarily, it can be noted that the midrib of almost all sorghum leaves lies on a plane that is perpendicular to the ground plane. In other words, from a top down view of the leaf, the midrib does not appear to curve at all. Furthermore, it is almost always the case that the point of inflection along a leaf's midrib is also the highest point on the leaf (or the point furthest from the ground plane). This would mean that a parabolic representation of a leaf would have an axis of symmetry very close to the vertical vector. For these two reasons, all parabolic leaf models are constrained to planes that are perpendicular to the ground plane while also having axes of symmetry equal to the vertical vector.

Parabolic RANSAC Leaf Detection Algorithm

1. The constrained plane in which all points within the cluster will be projected is found.
 - a. A basic linear RANSAC algorithm is run on the cluster with no angle constraint, a distance threshold of 1 cm, and a minimum model size of 20 points for a total of 10,000 iterations. This line is used to find a directional unit vector that roughly points along the length of the leaf.
 - b. The normal vector of the constrained plane is found using the cross product of the directional unit vector of the leaf and the vertical vector.
 - c. A point on the plane is found using the centroid of the inliers from the linear RANSAC in part 1a. Together, the point and normal vector fully define the constrained plane.
2. All of the points in the cluster are projected onto the surface of the constrained plane, and the associated Euclidean distances from each point to the plane are also stored.
3. A rotation matrix that rotates the constrained plane to be parallel with the XZ plane is found and the set of projected points are rotated to zero out the y-components for ease of future calculations.
4. Three points on the plane are randomly selected and the model parameters for a unique parabola are found given the vertical axis of symmetry constraint. The parabola must be concave down (to represent a physically accurate leaf) or else step 4 is repeated.
5. For all points on the plane, the Euclidean distance to the parabola is found. The total distance value for each point is equal to the projected distance from the original point to the plane plus the distance from the projected point to the parabola. If the total distance is less than the *distance threshold*, then the point is stored in a current list of inliers.
6. If the number of inliers is greater than the number of inliers from the previous best model as well as the specified *minimum model size*, then the current model is stored as the best model.
7. Repeat steps 4-6 until the *total number of iterations* has been reached.
8. The inliers from the best model are rotated back using the inverse rotation matrix found in step 3. The projected inliers form a clean parabola that resembles the leaf midrib while the reprojected inliers form all points along the detected leaf surface.
9. The algorithm is repeatedly run on the set of outliers from the previous result until there are no more leaves found in the cluster that meet the *minimum model size* constraint.

Table 4.2: Key Parameters in the Leaf Detection Algorithm

Number of Iterations	2,000
Distance Threshold	8 cm
Minimum Model Size	60 points

After running this parabolic RANSAC algorithm to completion on all of the clusters, a full model of all stalks and leaves in the original scene should now be extracted. This is the last stage of the leaf detection pipeline, however, the parabolic leaf models will now be used to detect optimal grasp points for the robotic manipulator.

4.2 Grasp Point Detection

4.2.1 Constraints Due to Limited Degrees of Freedom

In most robotic manipulation tasks, grasp point detection is based heavily on the geometry and pose of the target object. It is also the case that the kinematics of the robotic manipulator have a significant impact on the feasibility of detected grasp points. For example, a robotic arm with six or more degrees of freedom can essentially move its end effector to any position and orientation within the reach of the robot, making almost all detected grasp points in the scene feasible. However, given a manipulator with a limited number of degrees of freedom, a subset of seemingly ideal grasp points may be unachievable. In the case of the custom leaf grasping manipulator described in Chapter 3, only three degrees of freedom are used to move the leaf clamp end effector to a potential grasp point. The kinematics of this custom manipulator allow for the positioning of the end effector to any point in 3D space but the resulting orientation is uncontrollable and dictated by the system geometry, as shown in Figure 4.3 below. Therefore, the manipulator may position the leaf clamp at a grasp point, but the resulting angle of the clamp may make a successful grasp impractical or even impossible.




Figure 4.3: Uncontrollable end effector orientation at various grasp points

Given the consistent orientation and shape of stalks along the sorghum rows and the fact that the end effector could be applied to any point along the circumference of the stalk, the limited degrees of freedom of the manipulator was appropriate for stalk grasping. Conversely, for the task of leaf grasping

where the sorghum leaves have a wide variety of sizes, positions, and orientations throughout the field, the limited degrees of freedom heavily constrains the set of achievable grasp points.

4.2.2 Graspability Metric

Given the challenges and constraints from the limited degrees of freedom of the manipulator, it is convenient to develop a graspability metric that proactively ranks the feasibility of grasping detected leaf points. Therefore, the best grasp points can be extracted from a large set of leaf points based on the points that rank the highest in terms of graspability. This concept of graspability is derived from the work of Ruehl et al [27] where a similar metric is developed in grasping objects on planar surfaces. Naturally, the question then arises, what defines graspability in this context?

Here, *graspability* is defined as a measure of the quality of a leaf's pose as it relates to the kinematics and constraints of the manipulator. The remainder of this section is dedicated to creating a compact function, called the *graspability function*, that will produce a normalized score that reflects the estimated quality of a grasp at a specified point along the leaf. In order to develop this function, some standard notation must first be considered. In Figure 4.4 below, a sample leaf is shown along with various detected leaf points along its surface.




Figure 4.4: Leaf point notation for graspability metric

At each leaf point, r , three unit vectors are present: the surface normal, \hat{s}_r , the tangent vector, \hat{t}_r , and the plane normal, \hat{p}_r . When a leaf is detected from the parabolic RANSAC leaf detection algorithm described in Section 4.1.3, these three unit vectors are calculated for each point as follows. The plane normal, \hat{p}_r , has actually already been found in Step 1b in the parabolic RANSAC algorithm. The tangent vector, \hat{t}_r , is found by solving for the vector between the current point, r_k , and a neighboring point, r_{k+1} then normalizing, as shown in Equation 4.1.

$$\hat{t}_r = \frac{(r_{k+1} - r_k)}{\|r_{k+1} - r_k\|} \quad (\text{Eq. 4.1})$$

Lastly, the surface normal, \hat{s}_r , can then be found by taking the cross product between the tangent vector and the plane normal, as shown in Equation 4.2, to create an orthogonal set of axes centered at the leaf point.

$$\hat{s}_r = \hat{t}_r \times \hat{p}_r \quad (\text{Eq. 4.2})$$

Also shown in Figure 4.4 is the end effector vector, \hat{c}_r , a unit vector whose direction is parallel with the length of the last link of the manipulator when the end effector has reached the leaf point. The calculation of this vector comes from the inverse kinematics of the manipulator which are shown in Section 4.3.1. For now, however, it is enough to know the written definition given above.

Next, it is useful to create a binary function, g_r , which determines whether a leaf point is graspable given its surface normal. This is necessary as portions of a leaf with a large angle between the surface normal and vertical vector cannot be grasped by the end effector. In other words, the more vertical a part of the leaf becomes, the more difficult it is to grasp with the leaf gripper. This binary grasp function is shown below in Equation 4.3.

$$g_r(\hat{s}_r) = \begin{cases} 1 & \text{if } \cos^{-1}(\hat{s}_r \cdot z_g) < \theta_{max} \\ 0 & \text{otherwise} \end{cases} \quad (\text{Eq. 4.3})$$

Here, the θ_{max} parameter can be set to create a threshold on leaf points that are considered graspable based on the angle the surface normal makes with the vertical vector. During testing of the manipulator in the field, a value of $\pi/3$ or 60 degrees was used for θ_{max} . Figure 4.5 below provides a visualization of the raw point cloud shown in green as well as the surface normals of all detected leaf points. The surface normals are shown in yellow if they are considered graspable by the binary graspability function and red if they are considered not graspable.




Figure 4.5: Visualization of the binary grasp function (yellow = graspable, red = not graspable)

With all of the information presented above, a mathematical definition of graspability can now be created in the form of the graspability function shown in Equation 4.4 below.

$$G_r(\hat{s}_r, \hat{t}_r, \hat{p}_r, \hat{c}_r) = \left(\frac{1}{\alpha+\beta}\right) g_r(\hat{s}_r) (\alpha |\hat{c}_r \cdot \hat{p}_r| + \beta |\hat{s}_r \cdot z_g|); \quad \alpha, \beta \geq 0 \quad (\text{Eq. 4.4})$$

This function essentially outputs a normalized score that evaluates the quality of a grasp based on the angle between the end effector vector and the plane normal as well as the angle between the surface normal and the vertical vector. The first dot product term, $|\hat{c}_r \cdot \hat{p}_r|$, returns a value between zero and one where an output of zero represents a 90 degree difference between the end effector and the plane normal, and an output of one represents the case where the vectors are parallel, the best scenario. Here, α is known as the *yaw error gain* and varies the effect that this angle difference has on the overall output. Similarly, the second dot product term, $|\hat{s}_r \cdot z_g|$, returns a value between zero and one where an output of zero represents a 90 degree difference between the surface normal and the vertical vector, and an output of one represents the case where the vectors are parallel, the best scenario. β is known as the *roll error gain* and varies the effect that this angle difference has on the overall output.

The first term in the function, $\left(\frac{1}{\alpha+\beta}\right)$, normalizes the result so that the graspability score is between zero and one, where a score of one is considered best. The second term, $g_r(\hat{s}_r)$, is again the binary grasp function which will cause the entire equation to equal zero when a point is determined to be ungraspable. When a point is graspable, the binary grasp function has no effect on the score.

The two gain values can be adjusted depending on the application or end effector design, however, for in field testing it was often the case that α was set to 1 and β was set to zero. This is because the gripper showed a very strong ability to grasp leaves with up to 60 degrees of difference between the surface normal and vertical vector. Furthermore, it was determined that the angle between the leaf plane normal and end effector was the single most important variable by far. This angle difference is visualized depicted in Figure 4.6 below in order to provide some intuition on the graspability metric presented above.

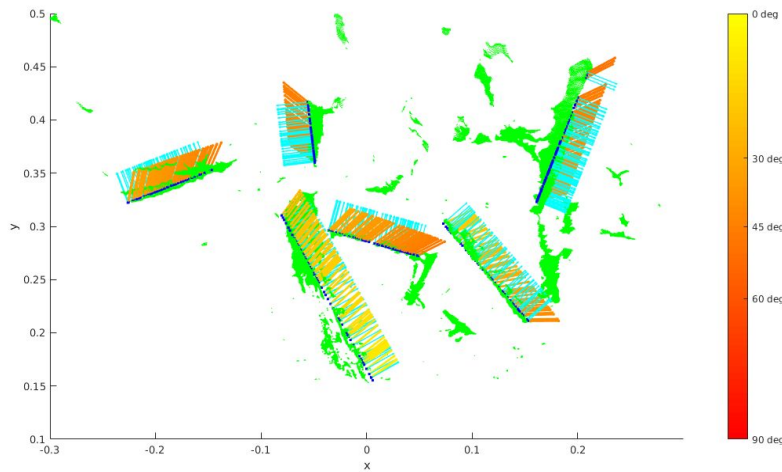


Figure 4.6: End effector vectors plotted over leaf plane normals

The plane normals are shown in light blue on all the detected leaf points whereas the calculated end effector vectors are shown in a color gradient between yellow and red. A yellow end effector vector represents a case where there is a zero, or close to zero, angle to the plane normal. Conversely, a red vector represents an almost 90 degree angle between the two vectors. As referenced previously, the end effector vectors are calculated from the inverse kinematics of the manipulator which are presented in the following section.

4.3 Manipulator Kinematics and Control

4.3.1 Inverse Kinematics

In order to determine the joint angles that will position the end effector to a given location in space, the inverse kinematics of the system must be solved. This is a crucial process as the robotic arm is controlled by passing in the desired joint angles to the controllers of the revolute joints. A simple model of a planar two link kinematic chain is shown in Figure 4.7 below.




Figure 4.7: Simple planar two link kinematic chain model

Given that the derivation of the inverse kinematics of such a system is so common in introductory robotics texts, the majority of the steps have been excluded and the final form is simply shown instead. The two joint angles, θ_1 and θ_2 , are calculated as shown in Equation 4.5 and 4.6 below.

$$\theta_2 = \pm \cos^{-1}\left(\frac{x^2 + y^2 - L_1^2 - L_2^2}{2L_1L_2}\right) \quad (\text{Eq. 4.5})$$

$$\theta_1 = \tan^{-1}\left(\frac{y}{x}\right) - \tan^{-1}\left(\frac{L_2 \sin \theta_2}{L_1 + L_2 \cos \theta_2}\right) \quad (\text{Eq. 4.6})$$

It should be noted that for most points, (x,y) , there are two configurations of the planar two link arm known as the elbow up and elbow down configurations. This fact is realized in the \pm sign in the expression for θ_2 . Depending on the sign of the second joint angle, the first joint angle will adapt given

that the equation for θ_1 is a function of θ_2 . Lastly, it should be noted that the *atan2* function, common in most programming languages, should be used in place of the inverse tangent function above when implementing these formulas in code. This will ensure that the inverse tangent properly returns angles within all four quadrants.

After these two joint angles have been found, the end effector vector referenced in Section 4.2.2 can be created using Equation 4.7 below.

$$\hat{c}_r = \langle \cos(\theta_1 + \theta_2), \sin(\theta_1 + \theta_2), 0 \rangle \quad (\text{Eq. 4.7})$$

A basic simulation was created to test the accuracy of the inverse kinematics. A visualization including the manipulator configuration for multiple grasp points along a leaf is shown in Figure 4.8 below. As before, the plane normals are displayed in light blue while the end effector vectors are shown in the yellow to red gradient.

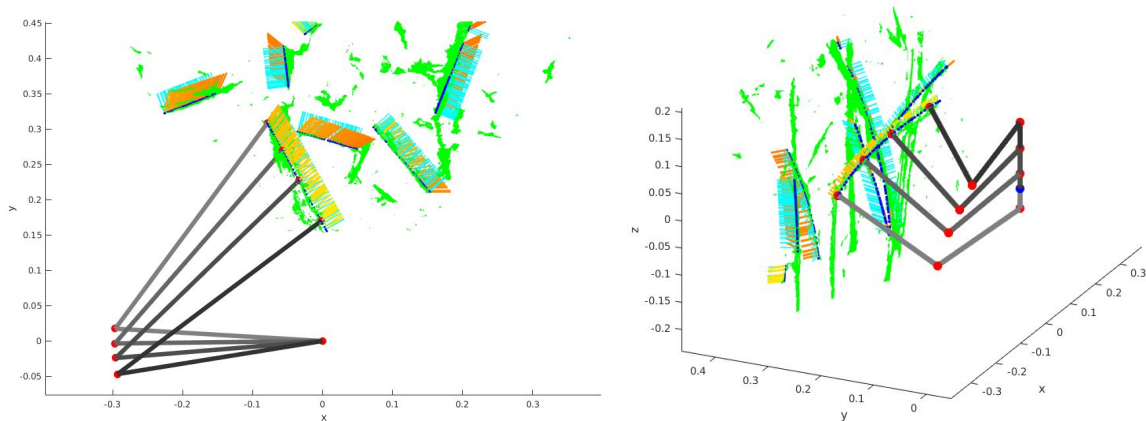


Figure 4.8: Manipulator configurations for multiple grasp points (left: top view, right: 3D view)

The inverse kinematics are solved online in C++ on an Intel NUC i7 processor in an average of about 50 microseconds. During the grasp point detection pipeline, the two joint angles are calculated for each of the potential grasp points in the scene. With an average of about 200 potential grasp points, the entire process of solving the inverse kinematics for all points takes about 10 milliseconds or one one-hundredth of a second.

4.3.2 Manipulator Control

The limited degrees of freedom that made grasp point selection so challenging has the significant advantage of making the control of the robot arm much simpler. Planning a trajectory for the manipulator is almost always as simple as parameterizing a short path with a few waypoints. The HEBI X-5 actuators achieve joint angle commands using an onboard PID controller. Therefore, the control of the manipulator requires solving the inverse kinematics for each location along the trajectory and feeding the joint angles into the PID controller. A simple global motion controller in ROS publishes joint state commands to the controller and continuously listens to the position feedback from the encoders on the motors. When the current sub goal has been reached, it publishes the joint angle commands for the next sub goal along the

trajectory. This process is run until the final goal point on the trajectory has been reached, in this case the grasp point on a leaf.

4.3.3 Trajectory Generation

When the infield process begins, the manipulator is first brought to the lowest position on the vertical post and the second link is rotated so it overlaps with the first link. This ensures that the arm assembly is not in the field of view of the Multisense stereo camera during imaging. The Multisense then records an image from the left and right frame and publishes the associated point cloud in ROS. At this point, the leaf detection pipeline is run and a specified number of optimal grasp points are selected. The manipulator is then placed in an “active” configuration. This active configuration is a set of joint angles that positions the end effector in the center of the row just past the shoulder joint. This configuration provides a sort of home position between grasp points and allows for simple, small joint commands to position the end effector throughout the scene. For each grasp point, a trajectory of four points is generated as shown in Figure 4.9.




Figure 4.9: Four step trajectory for leaf grasping

As with the trajectory for all grasp points, the first step is getting the arm in the active configuration. The second point along the trajectory, called the approach position, is 10 cm before the grasp point along the direction of the leaf plane normal. The third point, suitably called the leaf position, is the grasp point itself. Lastly, the fourth point, called the push position, is 5 cm past the grasp point along the direction of the leaf plane normal. This fourth and final position is added to ensure that the leaf is resting in the clamp of the end effector. Sometimes, the end effector does not quite make it all the way to the leaf, and pushing the clamp a little past the detected point is never the cause of an unsuccessful grasp. After the trajectory is completed, the clamp is closed, a spectral scan is taken, the clamp is opened, and the manipulator returns to the active configuration and awaits the next grasp point.

Chapter 5

Field Testing and Data Collection

5.1 Experimental Leaf Grasping

The accuracy of the leaf detection pipeline and the ability of the custom manipulator to grasp stalks was tested during the summer of 2018 in the grain sorghum breeding plots managed by Clemson University in Florence, SC. The manipulator was mounted to the vertical stage of the ground robotic platform and the robot was driven throughout the field and stopped randomly to collect image data from the Multisense and perform leaf grasping attempts. Although not used in this initial testing, the spectrometer system and fiber optics required for spectral data collection were also mounted to the robotic base and end effector. Images of the manipulator during this field testing session are shown below in Figure 5.1.



Figure 5.1: The leaf grasping manipulator during initial field testing

The minimum allowable graspability ranking of grasp points, as described in Equation 4.4, was set at 0.7, corresponding to a maximum allowable angle difference between the plane normal and end effector vector of about 45 degrees. For this, a set of optimal grasp points, capped at five, was passed to the global motion planner and trajectories were generated for grasping. After a simple performance test, the robotic arm succeeded to grasp 34 leaves out of 50 total attempts for an overall success rate of 68%. Later manual processing determined that all 50 of the detected grasp points belong to leaves recorded in

the original point cloud. Therefore, error was largely due to wind that forced the leaves away from their initial positions. A simulation of performance was built in RVIZ using the data recorded through ROS of the state of the robot as well as image and point cloud data. A snapshot of this simulation is shown in Figure 5.2 below.




Figure 5.2: RVIZ simulation of robotic leaf grasping

In the above figure, the small light green spheres correspond to the downsampled voxel grid produced during the initial stages of the leaf detection pipeline. Furthermore, an example of a detected grasp point is shown as the larger orange sphere which is more visible in Figure 5.3 below.




Figure 5.3: Close-up of detected grasp point in RVIZ simulation

5.2 Spectral Reflectance of Sorghum Stalks

5.2.1 Experimental Design

Initial testing of the reflectance spectrometer system was completed in the summer of 2017 in the sorghum breeding plots managed by Clemson University in Pendleton, SC. At this stage in the project the spectrometer system only consisted of a broadband light source and the SWIR sensor capable of measuring light intensity from 1350 - 2450 nm. The system setup was such that light source projected light through one end of a bifurcated fiber optic cable and out of a small probe attached to the end of the plant stalk manipulator (shown in Figure 3.2). The probe was designed to make direct contact with the stalk and measures its reflectance spectrum by rerouting the reflected light through the other end of the bifurcated fiber optic cable and into the input of the spectrometer. With these reflective spectra, the goal is to develop algorithms to predict compositional traits of the plants in an extremely fast and nondestructive manner.

Experimental design began by focusing on plants identified as ideal candidates for spectrometer readings based on biochemical, compositional diversity. This subset of sorghum consists of 48 plants from the bioenergy sorghum field. A total of 12 plots were selected throughout the field representing 6 genotypes (with genotype repetition to study potential field effects). One representative stalk was selected from each row in the plot, therefore, a total of 4 stalks from each of the 12 plots resulted in the 48 samples. Each stalk was scanned by the spectrometer system five times in various locations along the length of the stalk for a total of 240 spectra. This process was repeated during each monthly visit to the Pendleton facilities (the last weeks of July, August, and September). The stalks that are analyzed in the field are then cut and trimmed so that only a roughly two foot section of the bare stalk remains. These samples are weighed to obtain a wet weight and then bagged for drying and grinding to prepare for in-lab spectrometer ground truth readings.

5.2.2 Results and Analysis

The nature of the data collection provides the ability to analyze the data in a variety of resolutions. Viewing all 240 spectra at once generally provides too much data to visually comprehend. On the contrary, averaging all 40 scans for each genotype, as shown in Figure 5.4, provides the lowest resolution of data and provides a great way to view the differences between the average spectra of all genotypes at once.




Figure 5.4: Average reflectance spectra of various genotypes




Figure 5.5: Demonstration of variance in average reflectance spectra

A close up of the reflectance spectra of all genotypes between 1800 and 2050 nm is shown in Figure 5.5 above. This zoomed in view better illustrates the variation in spectra. Furthermore, this large dip in the reflectance intensity values is closely linked to the absorption of water. Therefore, insights like this can help to make predictions as to the overall moisture content of stalks.

In order to test the accuracy of a few of the hyperspectral vegetation index equations, values of moisture were predicted using the Normalized Difference Water Index (NDWI) from Gao et al. [33], as well as the Water Band Index (WBI) from Penuelas et al. [34]. The relation used for NDWI is shown in equation 5.1, while the relation for WBI is shown in equation 5.2 below.

$$NDWI = \frac{R_{860} - R_{1240}}{R_{860} + R_{1240}} \quad (\text{Eq. 5.1})$$

$$WBI = \frac{R_{900}}{R_{970}} \quad (\text{Eq. 5.2})$$

The results of the NDWI moisture equation are plotted against the true moisture values in the left side of Figure 5.6 while the results of the WBI moisture equation are shown on the right side.

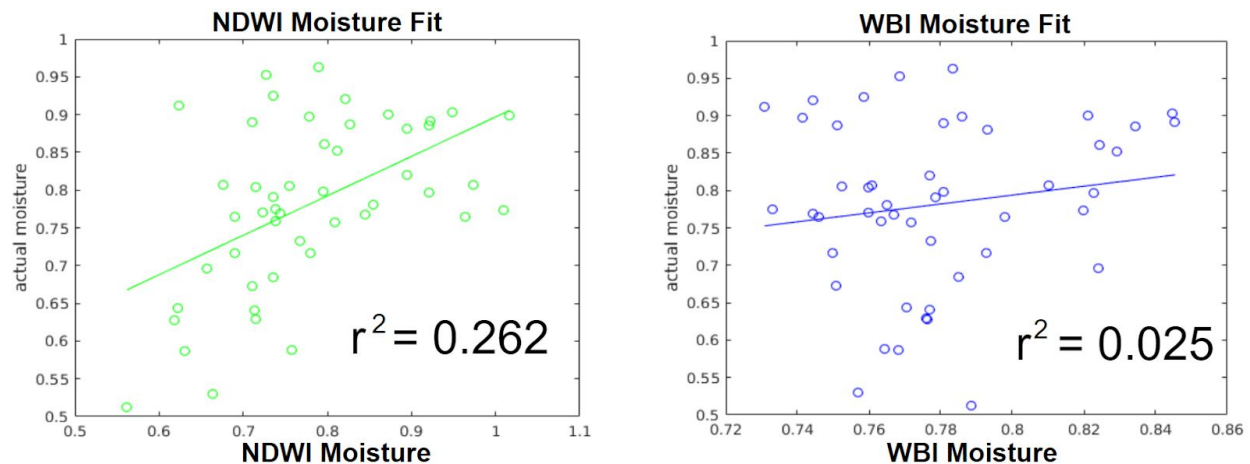


Figure 5.6: Results of using NDWI and WBI vegetation indices to predict moisture

Clearly, the correlations between the two metrics and the true values are very poor. This result certainly seems to reject the use of some common hyperspectral vegetation indices to accurately predict compositional traits of sorghum. However, it should be noted that most vegetation indices are developed

and optimized from data collected on a certain crop and may not apply to other crops well. However, further work must be done to check the validity of more hyperspectral vegetation indices or to even develop custom metrics specific to sorghum.

Additionally, a basic machine learning approach was also used in an attempt to more accurately predict useful biochemical traits. A single hidden layer neural network is created that takes in an input vector that fully describes the reflectance spectra. Therefore, the input is a 200 element vector containing the light intensity values at each of the 200 recorded wavelengths. The hidden layer consists of a single neuron that applies weights to all 200 input values then passes the weighted sum with a bias included to a sigmoid activation function. The network is trained in MATLAB using the `trainbr` function to create minimize a loss function that is a combination of squared errors and weights. This style of training while also minimizing weight terms is known as Bayesian regularization.

The five scans from each stalk are averaged to create a representative spectrum from each stalk. Along with the true measured traits, the average spectra are passed into the network and the weights are updated according through backpropagation. For each month, 38 of the total 48 spectra are used for training while 10 spectra are withheld to use for validation. Cross validation was performed by partitioning the data randomly into these groups a total of three times. The predicted values are plotted versus the input true values from the training and test stages as well as the combination of the two. This single layer neural network was trained using the moisture data and the results are shown in Figure 5.7 below. Similarly, the network was trained using starch data and the results are shown in Figure 5.8 below.

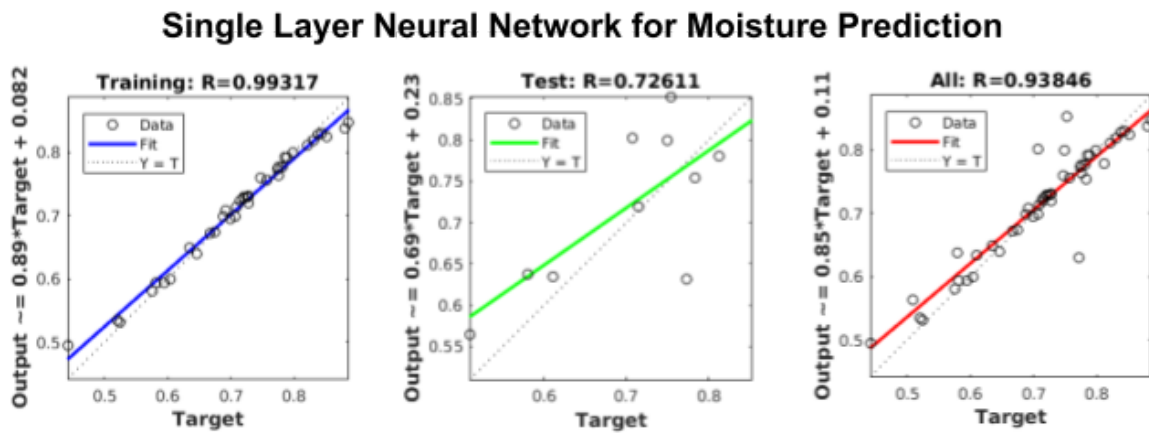


Figure 5.7: Correlations from training and test data using neural network for moisture prediction

Single Layer Neural Network for Starch Prediction




Figure 5.8: Correlations from training and test data using neural network for starch prediction

The above figures that portray the results from the single layer network show textbook signs of overfitting, and thus may not extrapolate well to larger data sets. The training error for both moisture and starch prediction is extremely low while the test error is much higher in comparison. This is most likely due to the fact that the size of the training data set is so small, much smaller than the number of features in the model. The development and use of this model is not meant to develop a robust model, however it does work to explore the potential of machine learning for this application. It is clear that future work must be applied to this area in order to explore the potential for learning based methods of biochemical trait prediction.

5.3 Spectral Reflectance of Leaves

5.3.1 Experimental Design

Initial testing of the full reflectance spectrometer system (all three spectrometers included for a combined reflectance band of 450 - 2450 nm) was completed in the summer of 2018 in the sorghum breeding plots managed by Clemson University in Florence, SC. A simple experiment was created to test the feasibility of predicting chlorophyll content in leaves, mainly sorghum leaves. Therefore, a sample size of 50 leaves was collected that consisted of 35 healthy grain sorghum leaves as well as 15 corn leaves of varying moisture/dryness levels. These corn leaves were added to the experiment in order to add a wide range of variance in chlorophyll values.

All 50 leaves were automatically grasped using the leaf grasping manipulator a total of 3 times each in various positions along the leaf. Therefore, 150 reflectance spectra were collected in total. The location of each automated scan was marked lightly on the leaf and a ground truth manual measurement was made in the same location using an Apogee MC-100 Chlorophyll sensor.

5.3.2 Results and Analysis

The reflectance spectra from all 150 leaves is plotted are plotted below in Figure 5.9. It can be seen that there is a wide range of variation between the minimum and maximum reflection percentage

values for each scan. This is largely due to the variance in leaf position and slight change in distance between the leaf's surface and fiber optic cables between each scan.

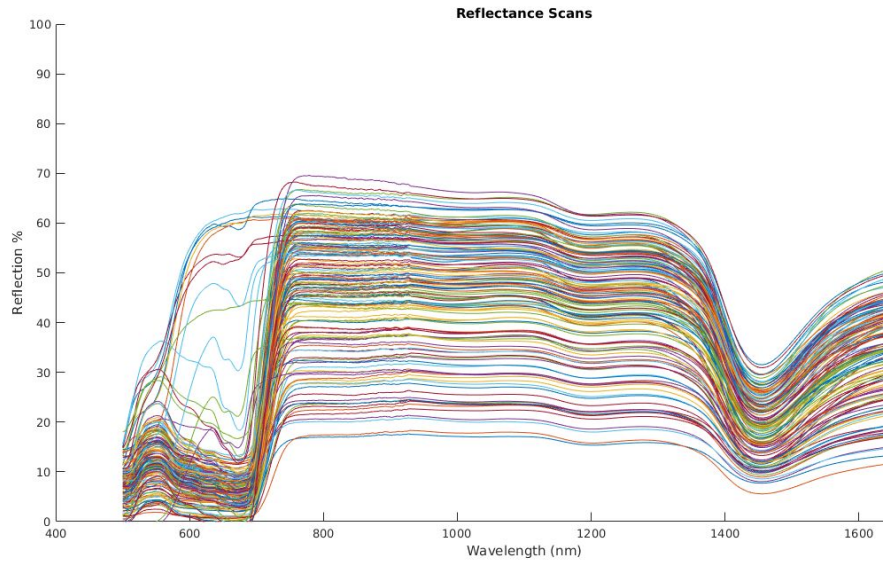


Figure 5.9: Raw reflectance spectra of leaves from chlorophyll test

However, this variance can be removed through a simple process of normalizing all reflectance scans. This is done by essentially reweighting the magnitude of each spectrum based on its raw minimum and maximum values so that each spectrum is contained between a normalized value of 0 and 1. The equation used to perform this normalization is shown in equation 5.3 below. The resulting normalized reflectance scans are shown below in Figure 5.10.

$$Normalized\ Reflectance = (max - min) \left[[raw\ spectrum] - \left(\frac{max-min}{2} + min \right) + \frac{1}{2} \right] \quad (Eq. 5.3)$$

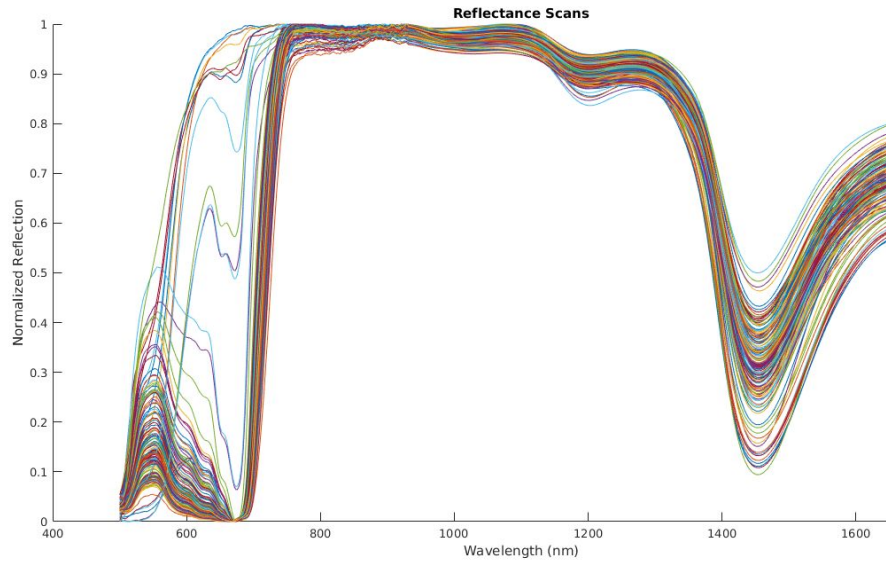


Figure 5.10: Normalized reflectance spectra of leaves from chlorophyll test

These normalized spectra were then used as input vectors to a similar single hidden layer neural network as described above in Section 5.2.2. A hidden layer with a single neuron applies weights to each of the ~ 1500 reflectance values for each wavelength in the spectrum. The neuron then adds a bias term and passes the sum to a sigmoid activation function. The output is then considered the predicted chlorophyll content. The network is trained using 85% of the 150 data points as training data while 15% of the data was reserved for validation data. Cross validation was performed by partitioning the data randomly into these groups a total of three times. The predicted values are plotted versus the input true values from the training and test stages as well as the combination of the two. The predicted values are plotted versus the ground truth values from the training and test stages as well as the combination of the two in Figure 5.11 below.




Figure 5.11: Correlations from training and test data using neural network for chlorophyll prediction

The above figure shows very promising results for the system's ability to predict chlorophyll content in leaves. The test error is slightly lower than the training error and achieves a rather high correlation coefficient of 0.93. The effects of potential overfitting here are minimal most likely due to the presence of a more defined trend, but also because the size of the training data was significantly larger. In general, it seems as though a broader reflectance spectrum is more effective at predicting compositional traits, especially chlorophyll, in leaves compared to the initial testing performed on sorghum stalks.

Chapter 6

Conclusion

This thesis detailed the process of developing a rapid and robust method for grasping sorghum leaves in the field as an automated method of performing reflectance spectroscopy on leaves. The development of a compact, in-field spectrometer system for high throughput spectral measurements was presented in detail. Additionally, the design of a custom manipulator and end effector specialized for the task of leaf grasping was explored in full. A novel leaf detection pipeline was developed in order to segment leaves from the surrounding environment in highly noisy point clouds. With these isolated leaf models, a graspability metric was applied to all leaf points in an effort to select the most optimal grasp points. A relatively simple but function strategy was used to develop trajectories for the arm to follow in the process of leaf grasping. Ultimately, the manipulator was able to grasp 34 leaves out of 50 total attempts for an overall success rate of 68% despite the presence of wind and unconstrained field effects. Lastly, an overview of in-field testing and experimental design for hyperspectral stalk analysis was given as well as a few methods used to predict compositional traits. Although initial work to develop accurate models achieved average performance, the start of more complex algorithms is explored with promising results.

Therefore, future work includes exploration into machine learning algorithms used to predict compositional traits of sorghum from reflectance spectra. To realize this goal, a much more extensive set of training data is needed to prevent overfitting and ultimately lead to low test error in practice. The vast advantages of nondestructive hyperspectral analysis of crops make this pursuit a worthwhile one. In order to improve the grasp success rate of the manipulator, more thorough forms of control can be applied. For example, visual servoing techniques are commonly used in manipulation tasks where the target object is unconstrained in some way. Therefore, actively tracking the position of a detected grasp point as the leaf moves throughout space could help create a more reliable system. Similarly, the geometric methods for leaf segmentation were suitable and reliable here, but state of the art research in agricultural perception often includes the use of some form of deep learning algorithms. Another approach to improve the detection and grasping pipeline would be to develop and train a network to segment leaves from 2D images and then associate grasp points to 3D points using depth data in post. These methods may even work to improve the functionality of the system by detecting more leaves despite the presence of occlusion and overlapping in the scene. Regardless, the foundational work presented in this thesis to develop a high throughput method of leaf grasping and contact sensing will hopefully make these future pursuits simpler and more accessible.

References

- [1] “The Future of Food and Agriculture - Trends and Challenges.” *Food and Agriculture Organization of the United Nations*, 2017, www.fao.org/3/a-i6583e.pdf.
- [2] N. Alexandratos and J. Bruinsma, “World agriculture towards 2030/2050: The 2012 revision,” *Food and Agriculture Organization of the United Nations*, 2012.
- [3] “Monthly Energy Review.” U.S. Energy Information Administration, 26 July 2010, www.eia.gov/totalenergy/data/monthly/.
- [4] Regassa, Teshome H., and Charles S. Wortmann. “Sweet Sorghum as a Bioenergy Crop: Literature Review.” *Biomass and Bioenergy*, vol. 64, 2014, pp. 348–355.
- [5] D. Rosenow and J. Dahlberg, *Sorghum: Origin, history, technology, and production*. New York: John Wiley & Sons, 2000
- [6] Flavell, Richard. “From Genomics to Crop Breeding.” *Nature Biotechnology*, vol. 28, 1 Feb. 2010, pp. 144–145.
- [7] R. T. Furbank and M. Tester, “Phenomics technologies to relieve the phenotyping bottleneck,” *Trends in Plant Science*, Vol. 16, No. 12, pp. 635–644, 2011.
- [8] “Transportation Energy Resources from Renewable Agriculture.” ARPA-E, 18 June 2015, arpa-e.energy.gov/?q=arpa-e-programs/terra.
- [9] Roberts, Dar, et al. “Hyperspectral Vegetation Indices.” *Hyperspectral Remote Sensing of Vegetation*, 2011, pp. 309–328., doi:10.1201/b11222-20.
- [10] Rouse, J. W., Haas, R. H., Schell, J. A., and Deering, D. W., Monitoring vegetation systems in the great plains with ERTS, in *Third ERTS Symposium*, NASA SP-351, NASA, Washington, DC, Vol. 1, p. 309-317. 1973
- [11] Curran, Paul, “Reflectance spectroscopy of fresh whole leaves for the estimation of chemical concentration,” *Remote Sensing of Environment*, vol. 39, issue. 2, pp. 153–166, February 1992.
- [12] “Field Scanalyzer - Automated Outdoor Phenotyping.” LemnaTec, www.lemnatec.com/products/field-phenotyping/field-scanalyzer/.

- [13] Bac, C. et al, “Harvesting robots for high-value crops: State-of-the-art review and challenges ahead,” *Journal of Field Robotics*, vol. 31, no. 6, pp. 888–911, November 2014.
- [14] Tanigaki, Kanae, et al. “Cherry-Harvesting Robot.” *Computers and Electronics in Agriculture*, Aug. 2008, www.sciencedirect.com/science/article/pii/S0168169908000458.
- [15] Z. De-An, L. Jidong, J. Wei, Z. Ying, and C. Yu, “Design and control of an apple harvesting robot,” *Biosystems Engineering*, vol. 110, pp. 112–122, 2011.
- [16] E. V. Henten, J. Hemming, B. van Tuijl, J. Kornet, J. Meuleman, J. Bontsema, and E. van Os, “An autonomous robot for harvesting cucumbers in greenhouses,” *Autonomous Robots*, vol. 13, p. 241258, 2002.
- [17] Christopher, Lehnert, et al. “Autonomous Sweet Pepper Harvesting for Protected Cropping Systems.” *IEEE Robotics and Automation Letters*, vol. 2, no. 2, Apr. 2017, pp. 872–879.
- [18] Nuske, S.; Achar, S.; Bates, T.; Narasimhan, S.; Singh, S. Yield estimation in vineyards by visual grape detection. *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, San Francisco, CA, USA, 25–30 September 2011; pp. 2352–2358.
- [19] Cerutti, Guillaume, et al. “A Parametric Active Polygon for Leaf Segmentation and Shape Estimation.” *Advances in Visual Computing Lecture Notes in Computer Science*, 2011, pp. 202–213., doi:10.1007/978-3-642-24028-7_19.
- [20] Xia, Chunlei, et al. “Plant Leaf Detection Using Modified Active Shape Models.” *Biosystems Engineering*, vol. 116, no. 1, 2013, pp. 23–35., doi:10.1016/j.biosystemseng.2013.06.003.
- [21] Jingwei Guo, Lihong Xu. "Automatic Segmentation for Plant Leaves via Multiview Stereo Reconstruction", *Mathematical Problems in Engineering*, Volume 2017 (2017), Article ID 9845815
- [22] H. Baweja, T. Parhar, O. Mirbod, and S. Nuske, “A deep learning pipeline for high-throughput measurement of plant stalk count and stalk width,” *Conference on Field and Service Robotics*, Zurich, Switzerland, September 2017.
- [23] Ahlin, Konrad, et al. “Autonomous Leaf Picking Using Deep Learning and Visual-Servoing.” *IFAC-PapersOnLine*, vol. 49, no. 16, 2016, pp. 177–183., doi:10.1016/j.ifacol.2016.10.033.
- [24] Mueller-Sim, Tim, et al. “The Robotanist: A Ground-Based Agricultural Robot for High-Throughput Crop Phenotyping.” *2017 IEEE International Conference on Robotics and Automation (ICRA)*, 2017, doi:10.1109/icra.2017.7989418.
- [25] Mueller-Sim, Tim. “Development of a Ground-Based Robot for High-Throughput Plant Phenotyping.” *Carnegie Mellon University*, 2017.

- [26] Jenkins, Merritt. “Detecting and Grasping Sorghum Stalks in Outdoor Occluded Environments.” Carnegie Mellon University, 2017.
- [27] Ruehl, Steffen W., et al. “Graspability: A Description of Work Surfaces for Planning of Robot Manipulation Sequences.” 2011 IEEE International Conference on Robotics and Automation, 2011, doi:10.1109/icra.2011.5979779.
- [28] Graybeal, Jack, et al. “Spectroscopy.” Britannica, www.britannica.com/science/spectroscopy.
- [29] “Broadband Infrared Tungsten Bulb.” ThorLabs, www.thorlabs.com/newgrouppage9.cfm?objectgroup_id=8554&pn=HEP3965.
- [30] “Fiber Optic Coupling.” Newport, www.newport.com/n/fiber-optic-coupling.
- [31] “Newport guide to light collection through optical systems.” Newport, <https://www.newport.com/t/light-collection-and-systems-throughput>.
- [32] Yendrek, Craig, et al. “High-Throughput Phenotyping of Maize Leaf Physiological and Biochemical Traits Using Hyperspectral Reflectance”, American Society of Plant Biologists, January 2017
- [33] Gao, B. NDWI: A normalized difference water index for remote sensing of vegetation liquid water from space. *Remote Sensing of Environment*. 58:257–266, 1996.
- [34] Penuelas, J., Pinol, J., Ogaya, R., and Lilella. I., *Estimation of plant water content by the reflectance water index WI (R900/R970)*. *International Journal of Remote Sensing*. 18:2869-2875, 1997.