

# Unmanned Aerial System for Monitoring Crop Status

Donald Ray Rogers III

Thesis submitted to the Faculty of the  
Virginia Polytechnic Institute and State University  
in partial fulfillment of the requirements for the degree of

Master of Science  
in  
Mechanical Engineering

Kevin B. Kochersberger, Chair

John P. Bird

Jerzy Nowak

November 20, 2013

Blacksburg, Virginia

Keywords: UAS, Crop Monitoring, Precision Agriculture, Multispectral, NDVI

Copyright 2013, Donald R. Rogers III

# Unmanned Aerial System for Monitoring Crop Status

Donald R. Rogers III

## ABSTRACT

As the cost of unmanned aerial systems (UAS) and their sensing payloads decrease the practical applications for such systems have begun expanding rapidly. Couple the decreased cost of UAS with the need for increased crop yields under minimal applications of agrochemicals, and the immense potential for UAS in commercial agriculture becomes immediately apparent. What the agriculture community needs is a cost effective method for the field-wide monitoring of crops in order to determine the precise application of fertilizers and pesticides to reduce their use and prevent environmental pollution. To that end, this thesis presents an unmanned aerial system aimed at monitoring a crop's status.

The system presented uses a Yamaha RMAX unmanned helicopter, operated by Virginia Tech's Unmanned Systems Lab (USL), as the base platform. Integrated with helicopter is a dual-band multispectral camera that simultaneously captures images in the visible and near-infrared (NIR) spectrums. The UAS is flown over a quarter acre corn crop undergoing a fertilizer rate study of two hybrids. Images gathered by the camera are post-processed to form a Normalized Difference Vegetative Index (NDVI) image. The NDVI images are used to detect the most nutrient deficient corn of the study with a 5% margin of error. Average NDVI calculated from the images correlates well to measured grain yield and accurately identifies when one hybrid reaches its yield plateau. A secondary test flight over a late-season tobacco field illustrates the system's capabilities to identify blocks of highly stressed crops. Finally, a method for segmenting bleached tobacco leaves from green leaves is presented, and the segmentation results are able to provide a reasonable estimation of the bleached tobacco content per image.

# Acknowledgments

I would like to take this time to thank all of those who helped me along the way to achieving, what is in my opinion, my most significant academic work to date. Firstly, I would like to thank the members of Dr. John Bird's Mechatronics Lab for providing me with the hardware necessary to complete my project. I would also like to thank my committee members for answering all of my many questions and providing me with direction as I explored a topic completely new to the Unmanned Systems Lab. Specifically, I would like to thank Dr. Kochersberger for all of the opportunities he has given me. As an undergraduate sophomore with no experience, I'm not sure what Dr. K saw in me, but whatever it was he accepted me into the USL. Since then, he has pushed me to be a better student and engineer than I thought possible, and without him my graduate education would not have been possible.

Next, USLers past and present. To this day the lab holds my fondest memories at Virginia Tech. The lab was where I made my closest friends, learned many valuable skills, and gave me a place to just have fun and hang with the guys. It is to the credit of many past USLers, that I pursued a graduate degree, and in remaining in contact with them, gave me encouragement when I felt overwhelmed and there was no end in sight. Thank you in particular to those who took the immense time out of their schedules this summer to participate in the test flights I needed to complete this work. Thank you Gordon Christie for helping me get a handle on image processing and always taking the time to discuss a particular nuance with me. A special thank you goes out to my mentor, Ken Kroeger. Kenny was the first person to make me feel truly a part of the lab, and since then has been a constant source of advice, encouragement, and friendship. The USL and its members will always hold a special place in my heart, and I want to express my deepest gratitude for all the good times you've shown me.

To my family and friends I would also like to extend a thank you. I have the great fortune to remain in touch with my closest friends from my childhood, and through the phone calls, texts, and late night video game sessions you've helped me to have fun and enjoy my time spent away from academia. To my parents, thank you for your continued support and love. You've always helped me to think clearly and given me the courage to never settle for less than I am capable of. If not for you both, I wouldn't have grown to become the man I am today, and I hope that I will always continue to make you proud.

Lastly, I would like to thank my amazing fiancée, Tina. Without you I would not have kept my sanity through this arduous process. You are my constant partner and companion, you listen, comfort, and when necessary gave me a stern reminder never to quit. Thank you for sticking with me through the entirety of my college years and giving me a home to come back to every day. When it was looking bleak, and I began to question whether I should have pursued my master's you were there to remind why I went down this path, and that presented to me was an opportunity to better both of our futures. I love you with all my heart, and can't wait to be your husband.

# Contents

<b>Abstract</b>	<b>ii</b>
<b>Acknowledgments</b>	<b>iii</b>
<b>List of Figures</b>	<b>vi</b>
<b>List of Tables</b>	<b>viii</b>
<b>List of Algorithms</b>	<b>ix</b>
<b>Nomenclature</b>	<b>x</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Project Goals and Objectives.....	2
1.2 Organization.....	3
<b>2 Literature Review</b>	<b>4</b>
2.1 Remote Sensing in Agronomy.....	5
2.2 Crop Monitoring with Unmanned Aerial Systems.....	7
2.3 Remote Sensing of Corn Crops.....	9
2.4 Remote Sensing of Tobacco Crops.....	11
<b>3 Methods</b>	<b>14</b>
3.1 Multispectral Payload Development.....	15
3.1.1 Hardware.....	15
3.1.2 Software.....	21
3.2 Unmanned System Integration.....	23
3.3 Initial Test Site.....	27
3.4 Flight Mission Development.....	29
<b>4 Multispectral System Analysis</b>	<b>31</b>
4.1 Ground Images.....	32

4.1.1	Image Acquisition.....	32
4.1.2	Image Processing.....	33
4.1.3	Results.....	38
4.2	Flight Images.....	39
4.2.1	Image Acquisition.....	40
4.2.2	Image Processing.....	41
4.2.3	Exposure Test Results.....	48
4.2.4	N Rate Detectability Test Results.....	52
4.2.5	Grain Yield Comparisons.....	55
<b>5</b>	<b>Late-Season Tobacco Study</b>	<b>58</b>
5.1	Flight Mission Overview.....	59
5.2	Image Analysis.....	60
5.2.1	Image Acquisition.....	61
5.2.2	Row Space Estimation.....	62
5.2.3	Stress Estimation.....	67
5.2.4	Bleached Leaf Segmentation.....	71
<b>6</b>	<b>Summary &amp; Conclusions</b>	<b>76</b>
6.1	Summary of System Capabilities.....	77
6.2	Future Works.....	79
6.3	Concluding Remarks.....	81
	<b>Bibliography</b>	<b>82</b>

# List of Figures

Figure 3.1: JAI multispectral camera with lens. ....	16
Figure 3.2: The effect of the dichroic prism in the JAI camera. As light enters the camera it is split into Visible and NIR channels by the prism, this gives the camera its multispectral capability. ....	16
Figure 3.3: Spectral responses for the JAI multispectral camera.....	17
Figure 3.4: The multispectral payload hardware; JAI camera highlighted in green, trigger board highlighted in blue, and fit-PC highlighted in red.....	21
Figure 3.5: Front Panel of the multispectral system LabVIEW VI. ....	23
Figure 3.6: Yamaha RMAX UAV .....	25
Figure 3.7: The multispectral payload as integrated to the RMAX UAV. The mounting plate is painted maroon and is attached by rubber isolators to the payload tray underneath (gray aluminum). ....	26
Figure 3.8: Fertilizer application scheme for test field. Each subplot is labeled by its number followed by a letter for hybrid type and color-coded by its N rate. e.g. Subplot 203 is located in the 7 <sup>th</sup> range, is of hybrid type P1745HR and was sidedressed with 150 lbs/ac of nitrogen. ....	28
Figure 3.9: Detection of fertilizer response mission plan. ....	30
Figure 4.1: Ground Image captured on 6/12.....	33
Figure 4.2: Visible (left) and NIR (right) ground image of the same scene .....	34
Figure 4.3: NDVI image formed from Figure 4.2 .....	35
Figure 4.4: Two Stage Segmentation: a) initial NDVI image, b) GNDVI image, c) mask, d) segmented NDVI image.....	37
Figure 4.5: Visible (left) and NIR (right) image gathered during flight on 7/16.....	40
Figure 4.6: A single sample (red) has been selected from a visible flight image. ....	43
Figure 4.7: A visible flight image (left) and the resulting NDVI image (right). Notice how the shadow (red) in the visible image is no longer present in the NDVI image. ....	44

Figure 4.8: A bounding box (blue) drawn by the user during the sampling routine. The bounding box is drawn to identify one of the two subplots in the image.....	45
Figure 4.9: After bounding box selection random samples are chosen from each subplot. The first subplot and the corresponding samples are outlined in green, and the second subplot and its samples in red.....	46
Figure 4.10: Two samples taken from the same subplot in the same image. Sample a) has significantly more black pixels than sample b) and without proper consideration sample a) could appear to be from a less fertilized subplot. ....	47
Figure 4.11: Grain yield versus applied Nitrogen rate.....	56
Figure 4.12: Grain yield versus NDVI .....	57
Figure 5.1: Tobacco study mission plan .....	60
Figure 5.2: Sample tobacco study image.....	62
Figure 5.3: Tobacco image with user drawn bounding box (blue).....	64
Figure 5.4: Generating data ensemble; a) a single data vector, b) the entire data ensemble. ....	65
Figure 5.5: Single tobacco row spectrum .....	66
Figure 5.6: Averaged tobacco row spectrum .....	66
Figure 5.7: Tobacco image with row boundaries highlighted (red).....	67
Figure 5.8: Tobacco sample segmented by two methods; a) usual NDVI method, b) NDVI method with shadow rejection. Notice how the shadow in a) (circled in red) has been eliminated in b).....	68
Figure 5.9: Sample tobacco study NDVI image .....	70
Figure 5.10: Bleached leaf segmentation example: a) original visible image, b) green leaves segmented from visible image, c) bleached leaves segmented from visible image ..	74

# List of Tables

Table 3.1: JAI camera specifications .....	18
Table 3.2: Fit-PC fanless computer specifications .....	19
Table 4.1: Mean Pixel Intensity and NDVI for one sample ground image test.....	38
Table 4.2: Mean Pixel Intensity calculated by two methods .....	47
Table 4.3: ANOVA exposure test results. ....	49
Table 4.4: Connected Letters Report for exposure test.....	51
Table 4.5: Confidence Intervals for exposure test. ....	51
Table 4.6: ANOVA of P1184AM nitrogen rate detectability results. ....	52
Table 4.7: Connected Letters Report for P1184AM N rate detectability test.....	53
Table 4.8: Confidence Intervals for P1184AM N rate detectability test. ....	53
Table 4.9: ANOVA of P1745HR nitrogen rate detectability results. ....	54
Table 4.10: Connected Letters Report for P1745HR N rate detectability test.....	54
Table 4.11: Confidence Intervals for P1745HR N rate detectability test. ....	55
Table 5.1: Average NDVI for first 16 tobacco rows .....	70
Table 5.2: Moving average stress estimation.....	71
Table 5.3: Typical tobacco image content .....	72
Table 5.4: Bleached tobacco content estimation.....	75



# List of Algorithms

1	Form NDVI Image .....	35
2	Two Stage Segmentation .....	37
3	Sample Flight Image .....	48
4	Form NDVI Image with Shadow Rejection.....	69
5	Bleached Leaf Segmentation .....	73

# Nomenclature

<b>AGL</b>	Above Ground Level
<b>ANOVA</b>	Analysis of Variance
<b>COA</b>	Certificate of Authorization
<b>DFT</b>	Discrete Fourier Transform
<b>GNDVI</b>	Green Normalized Difference Vegetation Index
<b>LAI</b>	Leaf Area Index
<b>LiPo</b>	Lithium Polymer
<b>LSD</b>	Least Significant Difference
<b>MPI</b>	Mean Pixel Intensity
<b>N</b>	Nitrogen
<b>NDVI</b>	Normalized Difference Vegetation Index
<b>NIR</b>	Near-infrared
<b>PRI</b>	Photochemical Reflectance Index
<b>SRI</b>	Simple Ratio Index
<b>SSD</b>	Solid State Drive
<b>UAS</b>	Unmanned Aerial System
<b>UAV</b>	Unmanned Aerial Vehicle
<b>USL</b>	Unmanned Systems Lab
<b>VI</b>	LabVIEW Virtual Instrument
<b>X</b>	Horizontal image axis, refers to the width of an image
<b>Y</b>	Vertical image axis, refers to the length of an image

# Chapter 1

## Introduction

There is a growing interest in modern day farmers to find tools that allow them to make improved decisions regarding crop management. As global population continues to soar, the degradation of cultivated soil and land due to intensive agricultural practices has become a serious issue [1]. It is important now more than ever to maximize the crop yields in order to meet the growing demand of a hungry populace. However, many current agricultural practices seek to meet this yield demand by applying an overabundance of agrochemicals, which leads to soil erosion and significant chemical runoff. One way to increase yield, and simultaneously minimize fertilizer and pesticide applications, is to collect high resolution data of fields and intelligently apply agronomic inputs at the micro-level [2]. Methods exist today in which agronomists can estimate various crop parameters such as biomass and nitrogen concentration [3]. However, the flaw of these methods is that a small sample size is used to infer the status of an entire field or farm. To more effectively maximize yield while minimizing chemical applications, data needs to be collected field-wide and in resolutions that allow isolation of individual rows or single plants. In this way, the micro-level management of a crop becomes possible with the potential for optimizing the return of each individual plant. What is truly desired then, is a capability to monitor the status of an entire crop, down to individual plants, and report the relevant information that allows farmers to make use of new more intelligent practices.

To that end, this work suggest the use of an unmanned aerial system for acquiring field-wide multispectral imagery which can be quickly post-processed and the results translated in a way that provides farmers with a better understanding of their crop's status. The prime advantage of

using unmanned systems is that data can be collected from low altitudes, increasing resolution, and providing the ability to analyze individual plants. Such a system can fly over large fields relatively quickly and can repeat flight missions with much greater frequency than manned aircraft, even so far as to fly a mission multiple times a day. Data gathered from the flights can then be post-processed with little turnaround time and the most impactful results returned to farm managers. Unmanned crop monitoring is the tool that can help to usher in a new age of agricultural practices which sees farms managing their crops for increased yield with intelligent and responsible chemical applications.

## **1.1 Project Goals and Objectives**

The overall goal of this work is to investigate the potential capabilities of an unmanned aerial system for crop monitoring. Specifically, it was desired to analyze multispectral images and generate useful results that could be used for future crop status monitoring. Initial testing made use of a corn crop in which ground truth data regarding the nitrogen fertilizer levels applied to the crop was available, thanks to researchers in Virginia Tech's Crop & Soil Environmental Sciences department. The objective of this first round of testing was to assess the ability of the system to reliably detect the nutrient status of various subplots in the corn field. Further testing of the system occurred over a late-season tobacco field where the two desired results were overall stress estimation, and the isolation of green leaves from bleached leaves. A secondary objective was to verify the use of the Normalized Difference Vegetation Index for vegetation segmentation and improve upon current methods by means of simple machine vision techniques. One final objective was to develop an intelligent method for estimating crop row spacing. All of the above listed objectives were completed to some degree of success and illustrate just some of the capabilities of an unmanned aerial system for crop monitoring.

## 1.2 Organization

This work begins with a review of published literature, which covers past and present work in agricultural remote sensing, in Chapter 2. A brief history of remote sensing in agronomy is given, followed by examples of current research for crop monitoring. Special attention is given to the use of multispectral imaging as the preferred data set for crop parameter estimation. Also presented are examples of unmanned systems in use for crop monitoring that expand on the usefulness of those systems when compared to traditional means for data collection. The review closes with a look at the impact of fertilizer management techniques for the crops studied in this work. Chapter 3 focuses on the development of the multispectral system, from payload development to vehicle integration. An overview of the basic operation of the payload, as well as the method to communicate with that payload during flight, is given. Chapter 3 concludes by providing the framework for the initial tests of the unmanned system used in the crop monitoring project.

Chapters 4 and 5, which form the bulk of this work, contain the analysis of all the data gathered during the project. Chapter 4 presents the analysis of the images gathered over the initial test crop of corn. Images gathered on the ground are inspected first, followed by images collected during flight, in an effort to determine the potential of the multispectral system for field-wide monitoring. Segmentation methods used to isolate the vegetation in a scene are presented for both sets of images, as well as a routine to sample the flight images for statistical inference. In the final section of the chapter, results are given that describe the system's ability to detect nutrient deficient corn and estimate grain yield using the Normalized Difference Vegetation Index. Chapter 5 covers the final data set acquired for the project as part of a late-season tobacco study. A Fourier Transform method for detecting crop row spacing is detailed and applied for row by row stress analysis. The final work completed on the tobacco images is a segmentation method to isolate healthy green leaves from bleached leaves, and the potential for the method to estimate the concentration of bleached leaves to total vegetation. The last chapter of this work, Chapter 6, presents a conclusion of the crop monitoring capabilities of the developed unmanned system and gives recommendations for potential future work of the crop monitoring project.

# Chapter 2

## Literature Review

This chapter presents a review of literature previously published that covers the major topics of this work including: the application of remote sensing to agronomy, the integration of unmanned systems into crop monitoring solutions, and relevant research conducted on the two crops studied for the project. The review begins by looking at remote sensing as it applies to the fields of agronomy and crop science. A brief history of remote sensing beginning with satellite imagery and moving into low altitude manned aircraft is covered before specific modern works are discussed. Multispectral imaging is stressed as the preferred data collection method for crop monitoring and three different solutions for gathering multispectral data are detailed. To close the discussion on agronomic remote sensing the Normalized Difference Vegetative Index is described, as well as how modern researchers continue to use the index to generate useful results for a variety of applications. The second topic reviewed in this chapter is the use of unmanned systems for crop monitoring and remote sensing. Three publications are discussed, each of which stresses the advantage of an aerial platform that can gather high resolution imagery at low altitudes more frequently than other alternatives. The last two sections of the review cover the two crops studied for this work: corn and tobacco. Each section begins with a look at what topics researchers were interested in early on and what multispectral indexes proved to be most effective for estimating parameters like yield and chlorophyll concentration. Finally, fertilizer management techniques are discussed for both crops with a focus on the effectiveness of multispectral data for estimating the fertilizer level.

## 2.1 Remote Sensing in Agronomy

Agronomists and plant scientists have made use of remote sensing for decades in the fields of crop classification [4], yield prediction [5], and disease detection [6]. Remote sensing is a powerful tool because it offers a non-invasive method of gathering large bodies of data about a given crop. Two of the earliest and most important efforts in remote sensing for agriculture were initiated by the National Aeronautics and Space Administration (NASA) in the late 1960s and early 1970s. NASA partnered with Purdue University in 1966 to establish the Laboratory for Agricultural Remote Sensing (LARS) in 1966, intended to be a focal point for remote sensing in crop sciences. The LARS project began with research to determine the spectral reflectance of crop canopies using field spectrometers [7]. Later, aerial color and infrared photography was collected with manned aircraft for remote crop evaluation at the LARS center [8]. The other NASA project developed for terrestrial remote sensing was established with the launch of the Landsat-1 satellite in 1972 [9]. The first iteration of the Landsat carried an experimental multispectral scanner (MSS), which gathered data in the green, red, and two near-infrared (NIR) bands with an 80 meter resolution. Sensors aboard the Landsat continued to evolve over time and the current satellite, launched in 2013, collects data in 11 different spectral bands. Nine bands are used for imagery in the visible to shortwave infrared (SWIR) spectrum, including a pan band, each with a 30 meter resolution, and two more bands are used for thermal imaging with 100 meter resolution [10]. The Landsat program has been passed on to the United States Geological Survey (USGS), but NASA continues to fund many different remote sensing projects.

It is apparent from the NASA projects that a key to agricultural remote sensing lays in collecting data in multiple spectra. Spectral data can be gathered by various methods, the most popular being spectrometers, multispectral cameras, and more recently hyperspectral imagers.

Understandably, each device has its pros and cons. Spectrometers are used in conjunction with illumination sources to record the precise reflectance of narrow spectral bands. The ability to gather a multitude of narrow band reflectance data afforded by spectrometers lends them to a great many applications, such as detecting a brown plant-hopper infestation in rice crops [11]. However, in almost all cases spectrometers require calibration and additional signal processing to make the most of the collected data. Multispectral cameras, as the name suggests, take pictures

in multiple spectrums. Images provided by cameras offer human-readable data that can be easily understood before any additional processing. The downside to multispectral cameras is the bands covered in the images tend to be fairly large (200 – 500 nm), and consequently precise information of a crop's reflectance is unobtainable. Multispectral cameras are a useful tool as demonstrated by researchers at the University of California, Davis who used a multispectral camera acquiring images in four bands (blue, green, red, and NIR) to detect late blight in tomatoes [12]. Hyperspectral cameras seem to offer the best of spectrometers and multispectral cameras; they create spectral images with both spectral and spatial axes. The spectral resolution of hyperspectral cameras is much higher than multispectral cameras and can be on the order of just 10 nm. This allows for a continuous spectral reflectance to be extracted from each pixel in the image. The immense amount of data collected can then be processed with both imaging techniques and advanced data processing such as self-organizing maps and neural networks [13]. Hyperspectral cameras are very powerful tools, but as such they can be very costly and lie outside the budget of a university program or farm.

The prevalent use of multispectral data in agriculture led to the development of a tool which could simplify analysis and provide a quick means to identify and compare vegetation. The result was the formation of ratios of spectral reflectance, commonly referred to as a “vegetation index.” Researches in the early 1970's pioneered the most well-known and widely used index, the Normalized Difference Vegetation Index (NDVI) [14]. The equation to calculate the NDVI ratio is as follows:

$$NDVI = \frac{NIR-Red}{NIR+Red} \quad (1.1)$$

where *Red* is the reflectance of light in the red band and *NIR* is the reflectance in the near-infrared band. The rationale behind the formation of the NDVI lies in a trend for green vegetation to absorb red light and strongly reflect near-infrared light [15], while other substances typically pictured with vegetation, such as soil or water, tend to have low reflectance in the NIR band. This allows NDVI to be used for detecting and segmenting vegetation from other elements present in an image, and as such NDVI was used early on to successfully identify regions of live green plant canopy from Landsat images. It was later shown that plant leaves with greater concentrations of chlorophyll also absorb greater amounts of red light while reflecting significant



near-infrared light [16]. Higher concentration of chlorophyll allows for more energy to be created during the photosynthetic process and is generally thought of as a sign of healthy vegetation. Therefore, NDVI can be used not only to detect vegetation in a scene but also as a measure for relative plant health. NDVI does suffer from a few limitations. Foremost is the broad definition of the spectral bands used in the calculation. The index can be calculated with any number of different bands, each of which might have a different width and center wavelength. Depending on the multispectral data used, NDVI may or may not produce the desired results, and frequently the bands chosen for use in the NDVI are tailored to the specific application. Another limitation is the sensitivity of NDVI to soil effects. The reflectance properties of soil can change depending on content and moisture level, thereby resulting in a non-stationary NDVI value. This can lead to significant soil noise and potentially misinterpreting soil as vegetation. To overcome the limitations several other vegetation indices have been developed and seen wide use, such as the Soil-Adjusted Vegetation Index [17] and the Water Band Index [18]. However, NDVI is still frequently used for various applications such as estimating the Leaf Area Index (LAI) of a canopy [19], and estimating crop green biomass [20]. Overall, NDVI is a powerful tool, and with the proper planning and execution, can be used for vegetation segmentation as well as stress analysis.

## **2.2 Crop Monitoring with Unmanned Aerial Systems**

As this work has established, remote sensing for agricultural practices focuses on providing spatial and spectral data used for monitoring and analyzing a crop's status. In the past, satellite and manned aircraft photography have been successful in acquiring the necessary data, but there exists a demand today for higher resolution imagery collected more frequently and at a lower cost to farmers. Researchers at the Asian Institute of Technology (AIT) have proposed using unmanned aerial vehicles (UAV) as platforms for low altitude remote sensing [21]. The proposed system uses a rotary UAV to carry a variety of imaging payloads and combine the images with highly accurate GPS information for real-time field monitoring and mosaic map creation. Arguments are made that a combination of sensors and position information allows for a multi-dimensional approach at crop monitoring for lower cost and increased spatial and temporal

resolutions. The increase in spatial resolution is derived from the low altitude flights, but also from improved lightweight digital cameras, while the increase in temporal resolution is achieved by creating a system that is readily available to fly with short notice. Attention is given to the cost sensitivity of the proposed system with respect to developing areas that farm on a smaller and less commercial scale. The final statement made is that a low altitude remote sensing system can potentially provide vital crop and soil information while also providing a platform that could be applied to a variety of uses, such as disaster prediction and assessment, all at an affordable price for a developing region.

Several works have been published using systems similar to those proposed by the AIT. One such publication focused on using images gathered from an unmanned helicopter to create high resolution maps of corn and sugar beet fields [22]. The authors argue that helicopters are a strong choice for low altitude platforms due to their vertical takeoff and landing (VTOL) capabilities, as well their superior maneuverability at low speeds. The ability to launch and recover in a variety of locations makes helicopters much more useful to farmers whose land may not have sufficient space for fixed-wing aircraft to do the same. A low altitude system was also desired due to the infrequency in which satellite imagery can reach farmers during the traditionally rainy growing season. Development of a reliable method to reference images to global coordinates, essential for making accurate maps, was emphasized and the method applied to a roughly 0.5 acre corn field. Also presented was the segmentation of sugar beet rows from the surrounding soil, something not always possible with the low resolution data provided by manned aircraft. The segmented images were then used to create two maps of the sugar beet field; one processed to display LAI over the field and the other to display NDVI. The authors believe their “crop status maps” represent widely known vegetation indices that are useful for understanding the variability of crop growth. A final point made echoes what was said by the researchers at the AIT; unmanned systems provide a low cost solution for frequently collecting aerial imagery of a higher resolution than images acquired by satellite or manned aircraft.

Berni et al. [23] demonstrates the usefulness of UAVs for agricultural remote sensing focuses on monitoring vegetation with both narrowband multispectral and thermal bands. Once again, the authors argue that integrating remote sensors to UAVs will provide a low cost approach to meet the spatial, spectral, and temporal resolutions critical to successful agricultural monitoring. The

researchers used an unmanned helicopter to carry a multispectral and thermal imaging payload, and data was collected with multiple flights over an olive orchard, peach orchard and corn field. Much of the publication is dedicated to the proper calibration of the thermal imager and the creation of thermal orthomosaics for detecting water stressed trees in the peach orchard and corn field. Multispectral images were gathered in six narrow bands ranging from 490 to 800 nm. The authors chose to calculate three vegetation indices from the multispectral data: NDVI for LAI estimation, the Transformed Chlorophyll Absorption in Reflectance Index normalized by the Optimized Soil-Adjusted Vegetation Index (TCARI/OSAVI) for chlorophyll content estimation, and the Photochemical Reflectance Index (PRI) for potential water stress detection. Results from the thermal imager were compared to the calculated NDVI and PRI results to show that NDVI is a poor indicator of water stress, while on the other hand, PRI could be a strong indicator of water stress. LAI measurements were made on the ground with a commercial sensor and then the relationship between NDVI and LAI inspected for the olive orchard and corn field. NDVI correlated well with the ground data in both cases, but had a significantly stronger correlation ( $r^2 = 0.88$ ) with the olive trees. Finally, the olive tree chlorophyll content estimated by the TCARI/OSAVI index was strongly correlated with in-field measurements. Similar to Saguira's work, two maps were generated of a portion of the olive orchard; one was indexed to display chlorophyll content and the other to display LAI. All three of the selected vegetation indices showed significant results and proved that a UAV platform can be used successfully for a variety of agricultural monitoring purposes. In fact, a point is made that their low cost UAV remote sensing system yields results comparable, and possibly better than, those obtained from manned aircraft or satellite. The only limitation mentioned was the low flight endurance of the helicopter (20 min) and that if a similar system were to be made useful for large scale agriculture a fixed-wing UAV should be considered.

## 2.3 Remote Sensing of Corn Crops

Remote sensing has been used to study various factors of corn status for a number of years with an early focus on LAI and yield estimations. As expected, multispectral reflectance remains the prominent method for corn crop monitoring, model development, and parameter estimation. The

potential for multispectral data was initially explored to develop linear relationships between agronomic variables of interest and NDVI [24]. A later work by the same authors showed that NDVI is distinctively sensitive to soil present in the background of gathered images, establishing that models and estimates made with NDVI are most effective once the crop has reached peak LAI [25]. Predictive models for corn yield were later developed using NDVI and another vegetative index, the Green Normalized Difference Vegetation Index (GNDVI) [26]. Data was gathered by a digital camera mounted to a manned aircraft flying at 1500 m above sea level. Spectral bands collected were 30 nm wide in the Green and Red spectrum, and 200 nm wide in the NIR. By updating the models with data sampled over a period of months yield variation in the field can be explained. The authors also argue that including soil reflectance data in the models helps improve estimations as many soil factors, such as drainage and organic matter content, impact yield. That argument is counter to the widely accepted idea that minimizing soil noise is important for making accurate estimations of crop parameters.

Today, research on applying remote sensing for nitrogen stress and management has become increasingly popular as concern with nitrate leaching and water quality has grown. Nitrogen is the most limiting nutrient in corn and is critical to support the development of the crop. In the past between 50 – 80% of N fertilizer would be applied to a corn field prior to actual planting, with the remainder applied by sidedress after plant emergence. This leads to inefficient N uptake by the crop and considerable nitrate leaching. Modern N management schemes attempt to match N supply with demand of the crop during the growing season before reaching the reproductive stage [27]. The most effective schemes use split fertilizer applications with the majority of nitrogen applied as a sidedress after corn emergence, as opposed to before plating. It is desired then to relate some remotely sensed data to the nitrogen status of a corn crop. A technique was developed to create an “N reflectance index” by comparing the ratio of NIR (760 – 900 nm) to Green (520 – 600 nm) reflectance of corn given a particular nitrogen application to the same ratio of a the most fertilized corn [28]. The developed index was used for late season corn (V11 to R4 stage) and correlated well to normalized chlorophyll measurements and the total nitrogen concentration sampled from two different hybrids. Bausch states the greatest “need” for N status estimation begins around the V5 stage and would continue beyond tasseling and possibly to the R2 stage. However, during the early growing period soil background effects on canopy

reflectance presented a major obstacle decreasing the amount of useful data gathered. The publication concludes by insisting on an investigation to discover techniques which can be used to minimize soil noise and may allow for early season N status estimations.

Osborne et al. stated a few years later that the wavelengths for estimating important factors like N concentration should change throughout the growing stage [29]. Regression equations were created to predict the N concentration over a range of growth stages, and two of the regression equations were strongly correlated to the measured N content of leaves picked during the V5 – V7 and V6 – V8 stages. A hyperspectral camera with a narrow field of view ( $15^\circ$ ) was centered just 3 m over corn rows during data acquisition in an effort to reduce soil noise. The 1.4 nm spectral bands used in the early stage regression equations ranged from 420 to 980 nm, covering the entire visible and NIR spectrum. Although many bands were used to achieve high correlation the author states that the best predictions of N concentration were made with reflectance in the red and green spectra.

## 2.4 Remote Sensing of Tobacco Crops

Research topics on monitoring and managing tobacco mirror those of corn, namely a focus on estimating chlorophyll concentration, leaf area index, and nitrogen status. Recall that chlorophyll concentration is thought of as a good means for detecting the physiological state and stress conditions in vegetation. Multispectral reflectance data gathered from tobacco growing under controlled conditions gave rise to some initial vegetative indices used to determine chlorophyll concentration in leaves varying from dark green to yellow in color [30]. Based on those findings work was completed to verify the accuracy of multispectral reflectance for quantifying chlorophyll concentration [31]. It was shown that reflectance near 700 and 550 nm is sensitive to chlorophyll and as the concentration increases reflectance at both those wavelengths decrease. The response in reflectance at 700 and 550 nm correlates extremely well ( $r^2 = 0.98$ ) leading to use of both bands in simple vegetative indices with the reflectance at 750 nm, which was proven to be nearly independent of chlorophyll concentration. The two simple indices ( $R_{750}/R_{700}$  and  $R_{750}/R_{550}$ ) were compared with measured chlorophyll concentration and showed very strong correlations ( $r^2 = 0.95$  for both) in linear regression models. An argument is made to replace the

*Red* portion of the NDVI equation with a *Green* reflectance between 540 and 570 nm so that the resulting index should show increased sensitivity to chlorophyll concentration. When compared to the measured chlorophyll of the tobacco leaves the new index (label by the author as “green” NDVI) had a linear correlation coefficient of  $r^2 > 0.96$  and provided estimates with the smallest error seen in any trial.

As with corn, nitrogen nutrition is a limiting factor for tobacco growth and overall yield. Current research continues to stress nitrogen management schemes that apply fertilizer “as needed” over the duration of the growing season. The goal of “as needed” applications remains the same; improve the efficiency of application and reduce pollution caused by an overabundance of fertilizer. Critical to proper N management is the precise and timely monitoring of crop N status. A recent study was completed to investigate the change in tobacco canopy reflectance from seedling to maturation [32]. Also tested was the effect of different nitrogen levels on canopy reflectance. The results of the study showed that the reflectance changed significantly in different growth periods and with different fertilizer rates. As the tobacco matured reflectance in the NIR band (700 – 1000 nm) increased, while reflectance in the green and red bands (500 – 700 nm) decreased. However, once the plant has reached maturity the NIR reflectance drops severely to the lowest level seen in the study. The reflectance trends suggest that as the crop grows NDVI will increase steadily until maturation, at which point NDVI will decrease significantly, creating the potential for NDVI to detect the different stages of tobacco growth. The spectral characteristics of a tobacco canopy respond similarly to increased nitrogen fertilizer as they do to maturation. Most interestingly, is that when the crop was supplied with too much nitrogen there was a substantial decrease in reflectance in the NIR band. Again, the spectral trends lend credibility to NDVI as a method for detecting nitrogen status, with the caveat that a plant given with a surplus of fertilizer might have a similar NDVI as a plant that received a much smaller fertilizer application.

Just two years later, Svotwa et al. conducted an experiment to establish a relationship between two vegetation indexes and applied fertilizer rates for three different tobacco hybrids [33]. The indexes chosen were NDVI and the Simple Ratio Index (SRI), which is the ratio of NIR reflectance to Red reflectance. Both indexes were positively correlated with fertilizer rate applied to the tobacco plants, but the SRI edged out the NDVI for strongest correlation ( $r^2 = 0.91$  vs.  $r^2 =$

0.82). Although response to increased fertilizer varied among each tobacco hybrid the overall trend was for NDVI to increase significantly from no fertilizer applied to 50% applied, and for NDVI to peak off after 100% application. This trend is most easily seen in the logarithmic regression equation for NDVI to fertilizer rate. Although SRI outperformed NDVI in terms of overall fertilizer rate, NDVI had the stronger correlation to total nitrogen concentration when compared with SRI ( $r^2 = 0.91$  vs.  $r^2 = 0.80$ ). The publication concludes with an estimate that NDVI of 0.72 indicates optimum tobacco health, and that N management schemes using NDVI for nitrogen content estimation will be successful.

# Chapter 3

## Methods

In this chapter the development of the multispectral payload, the integration of the payload to an unmanned aerial vehicle, and the initial testing of the system are discussed. The hardware that makes up the system is covered first, with particular attention given to the chosen sensor and its unique capabilities. Following hardware is a brief look at the software used to control the payload. An overview of the basic operation of the payload software is given and the user interface used to control the software is presented. The second section of this chapter covers the integration of the payload to a UAV to complete the multispectral system. The specific vehicle used to carry the payload is discussed as well as the manner in which the payload is mounted to the vehicle. Closing the integration section is a brief overview of the air-to-ground radio link that allows for remote operation of the payload during flight. This chapter concludes by describing the setup of the initial test of the complete multispectral system. The physical parameters of the chosen test site are defined as well as its division into subplots and the application of the fertilizer rates. And finally the development of the initial flight mission, the methodology followed when executing the mission, and the considerations given to develop the mission plan are given.



## 3.1 Multispectral Payload Development

As stated in the title of this work it was desired to develop an entire unmanned aerial system capable of monitoring crop status. The process to develop a UAS begins with creating a payload based around an appropriately chosen sensor. Hardware is selected to power and operate the sensor and then software is written to control the payload. After bench testing, the payload can then be integrated to an existing UAV to create the final UAS. This section follows that same logic: select an appropriate sensor and the necessary hardware and then write software to control and monitor the status of the payload. The integration stage of the system development is described in the following section of this chapter. Emphasis is given to the sensor chosen for the project as it is the key for crop monitoring, and enough detail given to understand the step-by-step development of the entire multispectral system.

### 3.1.1 Hardware

Development of the multispectral payload began by choosing a sensor with the capability to monitor crop status. It was desired for the sensor to be a passive, visual sensor. The trend in autonomous systems has been to replace heavy, high power, active sensors with lightweight, low power, passive sensors [34]. As discussed previously, active sensors, like spectrometers, have been used for many years in the agriculture community to gather data about crop health. But as high quality cameras become cheaper and easier to use, visual sensors are beginning to see an emergence of use for agronomy as discussed in the previous chapter. The other requirement for the sensor was that it be able to gather data in multiple spectral bands. Specifically, the sensor needed to be able to collect data in the Red (600-700 nm) and NIR band (700-950 nm) in order to form the most basic and widely used vegetative index the NDVI.

The sensor selected for the project is the JAI AD-080 GE multispectral camera, which was already in the possession of the Virginia Tech Mechatronics Lab and is pictured below in Figure 3.1. The JAI has the ability to gather visible and NIR images through the same lenses using prism optics [35]. An illustration demonstrating the effect of the prism can be seen in Figure 3.2. As light enters the camera a dichroic prism allows for the precise separation of the visible and NIR spectrums [36]. Two separate CCDs are used to capture the light after passing through the

prism; the first uses a Bayer filter for color and the second, monochrome for the NIR. In this way, the JAI captures two images of the exact same scene, one in each spectrum. This is extremely advantageous for manipulating the images for analysis. Each spectrum is also captured with the same size CCD resulting in a visible and NIR image of the same resolution. With proper image acquisition there should be minimal misalignment of the two images and direct comparisons between the two spectrums can be made. These two features, simultaneous capture of both spectrums and resultant images of the same resolution, have a significant impact on forming the NDVI of the captured scene as discussed in the next chapter.



Figure 3.1: JAI multispectral camera with lens.

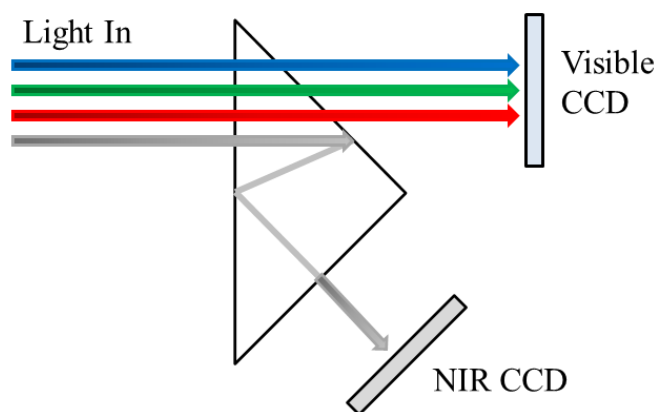


Figure 3.2: The effect of the dichroic prism in the JAI camera. As light enters the camera it is split into Visible and NIR channels by the prism, this gives the camera its multispectral capability.

The spectral response of the JAI multispectral camera is shown in Figure 3.3. The camera captures the full spectrum of visible light, and is sensitive to infrared light up to 1000 nm. It is evident then that JAI captures the sufficient spectral data needed to calculate the NDVI. Also, due to the way the visible light is captured a few other vegetative indices, such as the Green Normalized Difference Index, can be calculated as well. The specifics on how the vegetative indices are formed from images are discussed in the next chapter. It is important to note the peak sensitivity of both the Red and NIR spectrums as well as the width of each band. The peak sensitivity tells which wavelengths of light are best captured by the sensor, and the width tells how light is segmented into the four bands seen by the camera. A camera with narrower bands would be able to form more specialized vegetative indices such as the PRI [37]. Although the bands of the JAI are fairly wide, the multispectral camera is able to gather responses in the spectrums necessary to use calculate the NDVI.




Figure 3.3: Spectral responses for the JAI multispectral camera

Relevant specifications for the JAI camera are detailed in Table 3.1. Of import are its low weight, compact size, and low power consumption. Each of these is always a factor when selecting a sensor for a UAV payload. Less weight leads to longer flight times or additional

payload capacity, and low power means a reduced battery requirement, which in turn leads to longer flight times and so forth. It is good practice to design an aircraft payload with multiple platforms in mind. If a specific payload is intelligently designed and able to be flown on multiple vehicles it is likely to see increased use, and the added flexibility means that a single aircraft being out of order will not ground the payload. The size and weight of the JAI camera allow it to be put to use on several different aircraft, from large fixed-wing planes to small multi-rotor copters. On a final note, the specific model of the JAI camera chosen by the other researchers at Virginia Tech and used for this work features a GigE Vision interface. The GigE interface supports data rates up to 1000 Mbps over any 1000BASE-T Ethernet port. This provides ample data rate necessary to capture high quality images at frame rates up to 30 fps. Also, there exist several software packages with built in support for GigE Vision devices, making tasks like camera initialization and image acquisition simple and straightforward. All of the above mentioned specifications, along with its multispectral capabilities, make the JAI camera a natural choice for a commercial of-the-shelf sensor, and presents the best option for the project unless a custom sensor were to be developed.

Table 3.1: JAI camera specifications

<b>Specification</b>		<b>AD-080-GE</b>
Sensor	Visible NIR	1/3" Bayer mosaic CCD 1/3" Monochrome CCD
Full Frame Rate		30 fps
Resolution		1024 x 768
Control Interface		GigE Vision
Power		12 – 24 VDC, 7W
Dimensions (H x W x L)		55 x 55 x 98 mm
Weight		320 g

To meet the on-board computing need for the multispectral payload a fit-PC2 was selected. The fit-PC is a miniature fanless PC that runs Windows XP on an Intel Atom CPU. A dedicated operating system like Windows XP is extremely beneficial in that many commercially available software packages can be used for sensor control, air-to-ground communications, and data logging. With up to 1.6 GHz processing speed and 2 GB of RAM, the fit-PC is able to run the OS and several different software routines simultaneously without trouble. Critical to payload

integration is the fit-PC's use of a solid state hard drive (SSD) as the primary storage drive. As the multispectral payload was developed for unmanned aerial operations it is important that all hardware be able to withstand vibrations induced by the vehicle without failure, and the fit-PC's use of solid state storage removes the threat of read/write errors prompted by vibrations. Similar to the JAI camera, other important features of the fit-PC are its low power consumption, low weight, and compact size. At maximum load, the fit-PC only draws 9 W, making it very energy efficient for all the capability it affords the payload. In terms of connectivity, the computer includes a mini-SD port for additional data storage, four USB ports for peripherals and a gigabit Ethernet port for communications. The only drawback to the fit-PC is its tendency to run at high temperatures. The device is rated up to 70 °C, and heating was never an issue during bench top testing or unmanned operations, but in warmer climates a cooling solution would be advised. A table of relevant specifications for the fit-PC2 is shown below in Table 3.2, and more information on the device can be found at [38]. Overall, the fit-PC2 meets all the criteria for the on-board computing needs of the multispectral payload and was a natural choice.

Table 3.2: fit-PC fanless computer specifications

<b>Specification</b>		<b>Fit-PC2 Diskless</b>
Processor		Intel Atom Z530
Clock Speed		1.6 GHz
Memory		2 GB DDR2-533
Connections		4x USB 2.0
		1x Gbit Ethernet
		1x Mini-SD
Input Voltage		10 – 15 VDC
Power	Idle	4.5 W
	Load	9 W
Temperature	Idle	50 °C
	Load	70 °C
Dimensions (H x W x L)		27 x 115 x 101 mm
Weight		250 g

The next critical hardware component of the multispectral payload is a custom printed circuit board that handles the power distribution throughout the payload. The circuit board, referred to as the “trigger board”, was developed by the Mechatronics Lab to power several cameras for bench top testing. The board has four barrel connectors capable of outputting 12 V @ 1 A, and is

able to distribute all the power necessary to operate the multispectral payload. The board is called the trigger board because it also contains four lines for synchronous hardware camera triggers. Each trigger line was designed a simple flip-flop circuit and controls the frequency at which images are captured from the camera. The frame rate of the camera is set by the trigger line, and as was the case for this work, is hardware limited by the trigger board to 10 Hz. Coupled with the high resolution and wide field of view of the camera, a frame rate of 10 Hz is more than sufficient for the crop monitoring project.

A picture of the entire multispectral payload can be seen in Figure 3.4. Highlighted in the figure are the JAI camera, the fit-PC, and the trigger board. Also present is a gigabit Ethernet switch, GPS antenna, and lithium polymer (LiPo) battery. The switch is used to enable communication between the camera, the fit-PC, and the radio on board the UAV. In this way, a user on the ground can connect to the fit-PC over the vehicle radio and interact with the payload. Further discussion of the software used to control the payload as well as the air-to-ground communication link is described later in this chapter. The GPS antenna is a commercial USB enabled GPS unit and is used to log the position of the vehicle as each image is captured. Proper GPS logging is crucial so data gathered in the images can be referenced to ground truth data provided by the agronomists in charge of the field study. Power is provided to the payload by a 4 cell 5500 mAh LiPo battery. The battery is able to provide 12 V power for over an hour at full load. Lithium polymer batteries are a common choice of stable DC power due to their high energy-to-weight ratio and compact size. With all the hardware of the multispectral payload covered this chapter moves onto an overview of the payload software.



Figure 3.4: The multispectral payload hardware; JAI camera highlighted in green, trigger board highlighted in blue, and fit-PC highlighted in red.

### 3.1.2 Software

The software component of the multispectral payload is responsible for four main tasks: controlling the camera settings, capturing images, logging GPS coordinates, and remote operation. The first three tasks are all handled by the same commercially available software, National Instrument's LabVIEW. A single LabVIEW virtual instrument (VI) handles all aspects of the camera as well as the GPS log. Much of the camera control VI had already been developed by the Mechatronics Lab, only requiring some slight restructuring and pairing down to be used for the crop monitoring project. Many of the operations handled by the VI make use of code written by NI specifically for GigE Vision cameras, and included in the NI Vision Development Module. The fourth software task, remote operation, is handled by Windows Remote Desktop. Remote Desktop lets a Windows enabled PC control another such PC when connected over the same network. In this way, a computer on the ground can connect to the multispectral payload during flight and interact with the LabVIEW VI to control the camera. The air-to-ground link will be discussed in the next section of this chapter. What follows is an overview of the basic operation of the multispectral payload software.

1. **Connection** – A TCP/IP connection is established to the Visual and NIR cameras over the ports specified by the user in the Front Panel. Although the JAI exists physically as a single camera each CCD is addressed differently and appears as a separate camera in software, requiring two different ports and connections. A Serial connection is established to the GPS over the port specified by the user.
2. **Initialize Cameras** – A configuration file, stored in .xml format, is sent to each camera. The configuration file stores all the parameters used for controlling the camera such as: acquisition frame rate, auto exposure mode, and camera trigger.
3. **Monitor System** – A continuous video stream is established with each camera and is displayed on the Front Panel. This allows the user to monitor the scene in view of the cameras, and ensure proper performance of each camera. It is during this stage that the user can configure the cameras' exposure time and turn on/off the auto exposure. A change in exposure will briefly interrupt the video streams. Also during this stage the current position of the system is pulled from the GPS and displayed on the Front Panel. Note: No data is being stored in monitoring stage, only displayed to inform the user about the status of the system.
4. **Image Capture** – Video streams cease and triggered operation of the cameras begins. Each camera captures a frame with each pulse sent from the trigger board. The frames are written to the selected disk and labeled with the camera name and the current capture number. Each time a pair of frames is saved the current GPS position is written to a text file, headed by the current capture number. A user can start and stop image capture with a button on the Front Panel. When image capture is stopped the system returns to the monitoring stage.

Figure 3.5 shows the Front Panel of the multispectral payload LabVIEW VI. The Front Panel acts as the user interface to the software controlling the payload, and provides a multitude of information regarding the status of the system. The most important controls used for the basic operation of the system are detailed following Figure 3.5.



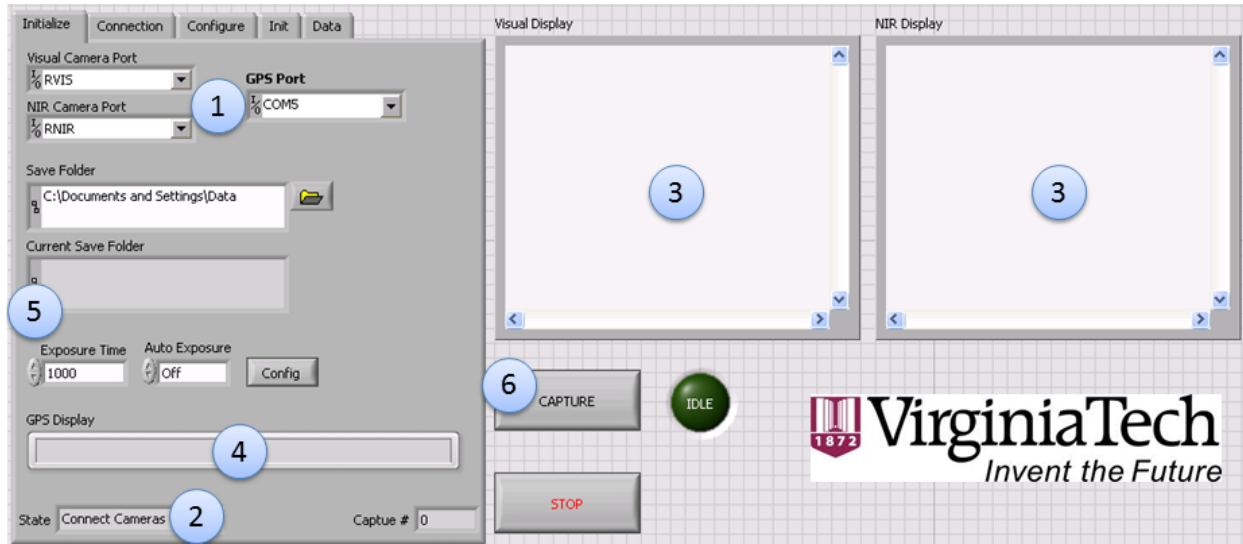


Figure 3.5: Front Panel of the multispectral system LabVIEW VI.

1. **Port Selection** – Controls for a user to select the port address for the Visible and NIR cameras, and the GPS antenna.
2. **Stage Indicator** – Displays the current stage the payload software is operating in.
3. **Image Displays** – These displays show either the video stream during the monitoring stage, or the most recent frame captured during the image capture stage.
4. **GPS Display** – Shows the current GPS position pulled from the GPS antenna.
5. **Configure Exposure** – Controls for setting the exposure time, and turning on/off the auto exposure.
6. **Capture Start/Stop** – A button that either starts or stops the image capturing software stage.

## 3.2 Unmanned System Integration

This section discusses the integration of the multispectral payload to a UAV including vehicle selection, hardware mounting, and basic unmanned operations. The baseline requirements for a vehicle to carry the payload are discussed first. Then, a brief description of the selected UAV is given, as well as the important components that allow for unmanned operations of the vehicle. Afterwards, the physical payload integration is covered, with a focus on simple effective mounting of the payload to the UAV. Finally, the air-to-ground link that allows for control of the

multispectral payload during flight is described in terms of hardware and remote operation during flight.

Given the size, weight, and power consumption of the multispectral payload, a set of baseline requirements can be established for a vehicle to carry the payload and successfully complete a crop monitoring flight. Firstly, the power consumption of the payload can potentially reach up to 28 W, assuming full draw on the fit-PC, camera, and Ethernet switch. To meet the power consumption requirement a LiPo battery has already been included in the payload design. While the inclusion of a power source allows the payload to operate independently of the vehicle, the battery adds significant weight to the payload. The size of the LiPo battery could be changed to reduce overall weight, but as it stands a suitable vehicle would need at least 5 lbs. of payload capacity, before factoring in mounting hardware. A more suitable requirement would be 8 – 10 lbs. of payload capacity. In this way, the system has flexibility regarding the mounting hardware used and should have spare capacity if additions to the payload are made such as additional sensors or batteries. To accommodate for the size of the JAI camera, when oriented in a Nadir view, the intended vehicle needs a minimum of 6.5 in of clearance between the payload bay and the ground. The clearance ensures that the camera lens does not come into contact with the ground before/after vehicle takeoff. This requirement might mean the addition of larger landing gear to smaller vehicles, or those that do not have large payload bays. Finally, the selected vehicle needs a minimum flight endurance of 10 minutes to allow for: takeoff, staging, at least a single pass over the quarter acre field, and landing. However, an endurance of 30 or more minutes would be recommended so more area could be covered without the need for multiple takeoffs/landings.

The UAV chosen as the platform for the crop monitoring project was the Yamaha RMAX helicopter. The RMAX has been operated by the USL for several years and was originally designed by Yamaha to provide assistance to aging farmers in Japan. The helicopter itself is over 6 feet in length from nose to tail and has a payload capacity of up to 50 lbs. A picture of the RMAX can be found in Figure 3.6. Integrated to the RMAX, and giving the vehicle its autonomous capabilities, is a wePilot flight controller. The wePilot is a commercially available autopilot that stabilizes the helicopter attitude and is capable of commanding constant altitude hover as well as waypoint navigation. When not in autonomous operation the RMAX is flown

like a traditional radio controlled (RC) aircraft by a highly trained pilot. RC control is an essential safety measure and is required by the Federal Aviation Administration for takeoffs and landings. Also onboard the helicopter is a Cobham COFDM radio. This radio is responsible for maintaining the command and control link between the ground station and the autopilot, and allows a user on the ground to command operation of the helicopter and monitor the vehicle's status. The Cobham also provides the air-to-ground radio link for payload control and operation through a separate port in the radio. A rotary UAV makes a great choice for crop monitoring due to its ability for slow controlled flight at low altitudes. This will afford ample time for any payload to gather high resolution data about the crops. However, a helicopter is also able to make high altitude, high velocity passes over a field, giving the system flexibility not always present with fixed wing aircraft. Overall, the RMAX was chosen as the vehicle platform for the crop monitoring project because of its reliability, large payload capacity, and flight mission flexibility.



Figure 3.6: Yamaha RMAX UAV

Once the multispectral payload was working reliably on the test bench, integration to the vehicle was rather straightforward. The primary concern with integrating the payload was mechanically fitting it to the helicopter. A payload tray had already been fabricated by members of the USL and featured 6 rubber vibration isolators that act as mounting points. The tray fits under the center of the helicopter and attached to the landing gear. A simple mounting plate made from medium density fiberboard was cut to fit the tray and mount to the isolators. The isolators allow for hardware attached to the plate to experience a reduced amount of vibration induced by the helicopter. Although some hardware was specifically chosen to be able to withstand vibrations, e.g. the SSD in the fit-PC, vibration isolation will assist with keeping the camera steady and

ensuring high quality imagery is taken. The camera is attached to the plate by a simple aluminum C-mount that allowed the camera lens to clear the payload tray without coming into contact with the ground. In this way, the landing gear of the helicopter is clear of the camera's field of view without risking damage to the lens. The C-mount is bolted to the tray and likewise the camera bolts to the C-mount. All of the other payload hardware is attached to the mounting plate by industrial strength Velcro. The Velcro keeps the hardware in place and securely attached to the plate during normal flight maneuvers, while allowing for relatively easy removal after a mission. Shown in Figure 3.7 is a picture of the multispectral payload mounted to the landing gear of the RMAX. It should be noted that rarely is it the case that hardware can be strapped to a vehicle and expected to function without fault. Several tests flight were conducted with earlier versions of the multispectral system, all of which failed, and it was only with more bench testing and redesign that the current multispectral system can operate successfully during flight.




Figure 3.7: The multispectral payload as integrated to the RMAX UAV. The mounting plate is painted maroon and is attached by rubber isolators to the payload tray underneath (gray aluminum).

As previously mentioned, a Cobham COFDM radio is used to communicate with the UAS during flight via an identical radio located inside the USL ground station. Connection to the air-to-ground link is accomplished by means of Gigabit Ethernet ports on both sides of the network. In the air, the multispectral system uses a GigE switch to connect the JAI camera and fit-PC to the radio on the helicopter. On the ground, another switch is used to connect a user's computer or

other device to the ground station radio. With a properly defined IP address, any device connected to the radios can communicate with any other device on the network. When powered on, the fit-PC allows for connections to itself with Windows Remote Desktop. Typical operation of the multispectral system involves establishing a Remote Desktop connection via the air-to-ground radio link to a PC in the ground station before takeoff. Once the connection has been established, the PC on the ground can monitor the payload during flight, configure the camera exposure time, and start/stop image capture. The Cobham radios operate at a frequency of 2.4 GHz and are able to maintain a strong link within one mile. The air-to-ground link is essential for unmanned operations as it allows users on the ground to control payloads in the air and ensure a successful mission.

### **3.3 Initial Test Site**

The goal of this section is to describe the site used for the initial testing of the multispectral system. The site chosen for initial testing was Virginia Tech's Kentland Farm, located in southwest VA. Kentland Farm is used for various research projects by the College of Agriculture and Life Sciences as well as 12 other VT departments. Unmanned operations flown by the USL have occurred at the farm since the lab's inception. It was discovered early in the summer of 2013 that Dr. Wade Thomason, of Virginia Tech's Crop & Soil Environmental Sciences department, was conducting a fertilizer rate study on a corn crop grown specifically for research. The fertilizer study was one of many tests conducted on the almost two acre corn field. After contacting Dr. Thomason, the subset of the crop dedicated to the fertilizer rate study was selected as the site for the unmanned test flight. A fertilizer study presents the perfect opportunity for the multispectral system to observe crops of varying status. As nitrogen fertilizer is the primary nutrient used on corn this gave the system an opportunity to detect stress in the form of nutrient deficiency. The controlled nature of the study provided sufficient ground truth data that can be used to determine the effectiveness of the multispectral system and support the results of the initial flight mission. The remainder of this section explains the work done by Dr. Thomason's team of researchers and the farmers at Kentland in organizing the field for the fertilizer rate study.

The goal of the fertilizer rate study devised by Dr. Thomason was to apply a range of nitrogen fertilizer amounts across the test field in order to estimate the optimum nitrogen uptake of each hybrid. To this end, four nitrogen rates were selected, beginning with the minimum of 50 lbs/ac and incrementing by 50 lbs/ac, until reaching a maximum fertilizer rate of 200 lbs/ac. An application scheme was developed in which each nitrogen rate was sidedressed to three subplots per hybrid, covering the entire test field. Figure 3.8 provides a visual representation of the application scheme applied to the test field and should serve as an aid for future discussion of nitrogen rates and subplot alignment. Exactly 1/3 of each rate (17 lbs/ac, 33 lbs/ac, etc.) was applied to the test field 5 weeks after planting, and the remaining 2/3 was applied after another 5 weeks. The corn was then allowed to mature until selected for harvest.

Range	Subplot		<b>Legend</b>
1	307 – A	308 – H	
2	305 – H	306 – A	50 lbs/ac
3	303 – A	304 – H	100 lbs/ac
4	301 – H	302 – A	150 lbs/ac
5	207 – A	208 – H	200 lbs/ac
6	205 – H	206 – A	<b>Hybrids:</b>
7	203 – H	204 – A	A – P1184AM
8	201 – A	202 – H	H – P1745HR
9	107 – A	108 – H	
10	105 – H	106 – A	
11	103 – A	104 – H	
12	101 – A	102 – H	

Figure 3.8: Fertilizer application scheme for test field. Each subplot is labeled by its number, followed by a letter for hybrid type, and color-coded by its N rate. e.g. Subplot 203 is located in the 7<sup>th</sup> range, is of hybrid type P1745HR, and was sidedressed with 150 lbs/ac of nitrogen.

The subset of the larger corn field used specifically as the test field for the fertilizer detection mission was only 30 feet wide, but stretches 360 feet deep. The field was planted on May 5<sup>th</sup> 2013 and growing in the test field specifically are two distinct corn hybrids: P1184AM and P1745HR. The test field was divided into 12 ranges; each range spans the entire width of the field and was 25 feet deep, with roughly 5 feet of bare ground between ranges. Every range was then sub-divided into 2 subplots, one of each hybrid. The complete subdivision of the test field makes for 24 subplots, each measuring 15 feet wide by 25 deep. Rows of corn planted in the

ranges are evenly spaced 30 inches apart, leading to 12 rows per range, and subsequently 6 rows per subplot. Because the row spacing was roughly even, there was not a clear division between subplots unlike there was with ranges. This necessitated the creation of a user guided tool to estimate the boundaries of the individual subplots during the analysis of images gathered from the flight mission. It should be noted that while every range contains a subplot of each corn hybrid, the planting of the hybrids does not follow a set pattern, i.e. hybrid P1184AM is not always the hybrid planted in the western subplot of the range, and neither is it always the eastern. Again, this required user input to identify the hybrid planted in a subplot during image analysis. Please refer back to Figure 3.8 to assist with visualizing the subdivision of the test field.

### **3.4 Flight Mission Development**

This final section presents the development of the first unmanned flight mission. The flight mission itself was a designed test of a UAS. Any time an UAS is to be used for data collection, or any other purpose, there exists significant risk to personnel, the unmanned vehicle, and any carried payloads. A mission plan was developed to keep the crew focused on the task at hand, thereby minimizing the above mentioned risks. Critical to successful flight missions are clearly defined goals and an efficient flight plan. The primary goal of the fertilizer detection mission was to gather multispectral images of the test field in order to detect the fertilizer level of each subplot. The stated goal made the objective of the flight mission clear and assisted in creating the flight plan. The flight plan was the set of waypoints used for autonomous navigation during the flight. Each waypoint had an associated GPS coordinate complete with a designated altitude above ground level (AGL). Also included with the flight plan was a desired speed to move from waypoint to waypoint, as well as any other flight maneuvers like turns, loiters, or changes in altitude. Other important aspects the mission plan were the position of ground station and selection of the takeoff/landing zone. The ground station should provide a clear view of the vehicle during takeoff/landing, and be able to maintain line of sight to the vehicle during at least one leg of the flight plan while remaining at a safe distance. Lastly, the landing/takeoff zone should be flat, free of debris, and well clear from any obstacles overhead.

For the detection of fertilizer response mission the flight plan was a simple straight line down the length of the test field. Waypoints were estimated at the center of each range, and the altitude set to 20 meters AGL, the lowest altitude for safe autonomous operation of the RMAX. Only a single pass at that altitude was needed to gather the full width of a given range. Figure 3.9 shows the flight plan overlaid in Google Earth (© 2013 Google). Please note the image pulled from Google Earth had not been updated in a couple of years and does not show the test field as was planted during this project. Two obstacles, the tree line to the left of the field and a set of power lines (labeled in black) running the length of the road, necessitated the use of a spotter in the field to warn the pilot if the vehicle was approaching too near to either obstacle. The ground station and landing/takeoff zone were positioned right of the field and roughly aligned with the 6<sup>th</sup> range. In this way, the ground station had a centered view of the field, and the helicopter can reach altitude before passing over the corn to reduce damage to the plants. The mission flight speed was set to 1 m/s, meaning a single pass over the field could be completed in less than five minutes. The developed flight plan was simple and focused on the efficient gathering of data. This would lead to the successful flight test, completed in July, which provided this work with its most important data set.




Figure 3.9: Detection of fertilizer response mission plan.



# Chapter 4

## Multispectral System Analysis

This chapter will focus on the analysis of the images gathered by the multispectral system for the fertilizer detection study. Images were primarily gathered on two separate occasions. The first section of this chapter discusses the initial images gathered on the ground. It was desired that an analysis of the payload and its potential be completed before integrated to any flight vehicle. In this way, the payload must demonstrate its value before risking a flight mission. The first section details the acquisition of these ground images, how they were processed, and their ability to detect the different nitrogen fertilizer rates. The other occasion at which data was gathered was the first successful flight of the fertilizer detection mission. The second section of this chapter discusses the flight images acquired during that mission, the method in which the flight images were processed for statistical analysis, and the results of that analysis. Particular attention is given to how vegetation in the flight images is segmented and then properly sampled as required for statistical inference. The chapter closes by comparing the results of the flight mission to those of Dr. Thomason's fertilizer rate study, in hopes that in the future the multispectral system can be used as a tool to aid in grain yield estimations.

## 4.1 Ground Images

It is common practice in the field of unmanned systems and engineering in general that a system undergoes various feasibility tests before a final implementation. This is especially important in unmanned systems, as every flight test risks damaging, or even worse loss of, the vehicle and any carried payloads. In this way, a potential payload must “prove” its effectiveness on the ground before being integrated to the vehicle. It should be noted that even the best designed and most thoroughly tested payloads may offer great results on the ground, but not necessarily translate those results to the air. As such, before any test flights could be completed images were taken by the payload as it was carried by hand over the test field. Recall, the goal of the multispectral system is to monitor crop status. The ability to detect a crop’s fertilizer or nutrient level is one way of completing the system’s goals. Therefore, for the payload to be considered feasible the ground images would need to show a potential to distinguish highly fertilized corn stalks from less fertilized ones.

### 4.1.1 Image Acquisition

The ground images were gathered on June 12<sup>th</sup>, 2013 when the corn was 5 weeks old. At this point most stalks were of the V4 – 5 stage and only 2.5 – 3 feet tall. The lighting condition was good and an exposure of 750  $\mu$ s was used for every image. This allowed the images to be taken by hand, as the payload could be held above single stalks. Figure 4.1 shows one such image gathered. The payload was walked down an interior row of the field, passing over every odd numbered subplot. Pauses were taken every few steps to ensure that at least three stalks of each subplot would be imaged with minimal blur. After the payload reached the end of the field handheld GPS measurements were made of the center of each subplot to create reference points for the GPS log created by the payload.



Figure 4.1: Ground Image captured on 6/12

Before analyzing the ground images, each image needed to be referenced to its corresponding subplot. This was done by correlating the GPS log created by the system to the handheld GPS measurements made in the field. Every time a picture is captured the log is updated with a new entry headed by the current number of images captured for easy cross-reference. In almost every case the handheld GPS measurements were a few feet off to the North West of the system log. This was easily accounted for and the images were sorted accordingly. Finally, each image was tagged by what subplot it was in and which nitrogen rate was applied.

#### 4.1.2 Image Processing

The first step to processing the ground images was to form the NDVI of each scene. Recall the general equation for NDVI can be represented as follows:

$$NDVI = \frac{NIR - Red}{NIR + Red} \quad (4.1)$$

Typically NDVI is calculated over a data vector gathered from a spectrometer. In that case, center wavelengths are chosen for the *NIR* and *Red* reflectances captured by the spectrometer. The most common center wavelengths are 675 nm and 905 nm for the *Red* and *NIR* responses

respectively [39]. Recall, the peak response wavelengths for the JAI multispectral camera are 610 nm for *Red* and 800 nm for *NIR*. Although the peak responses of the camera are not the typical center wavelengths used in NDVI calculations, the multispectral camera is able to capture sufficient spectral reflectance data to form the index. In the case of images taken by the multispectral system, NDVI is calculated over the entirety of each scene. The result then is not just a data vector, but a grayscale image with each pixel of the image relating to a value on the NDVI scale.

To form the NDVI of a scene recall that the multispectral system captures an image in both the visible and NIR spectrums at the same time through the same lens. Because both the visible and NIR images show the same scene, at the same time, pixel-by-pixel operations such as addition and division can be used to form the NDVI without any loss of data. It should also be noted that the visible and NIR images are the same resolution, and therefore neither image must be subsampled, which introduces blurring, before forming the NDVI. This preserves resolution and the clarity of the scene. If any of the camera properties mentioned above were different between the two images (timing, field of view, etc.) forming the NDVI could take significant estimations and photo manipulations, effectively reducing the quality of the data and comprising the ability to draw meaningful conclusions. Figure 4.2 demonstrates the above mentioned capabilities of the multispectral system.



Figure 4.2: Visible (left) and NIR (right) ground image of the same scene

Looking back to Equation 5.1, it is apparent that two separate images are needed to form the NDVI of a scene. The *NIR* data used in calculating NDVI is simply the NIR image gathered by

the multispectral system. On the other hand, the *Red* data must be separated from the visible image captured by the system. However, this process is already mostly completed for the user on the computer acquiring images from the camera. The sensor used to capture the visible scene is a 1/3" Bayer CCD, so called because it uses a Bayer color filter array. The Bayer filter allows visible light input to the sensor to be output in three distinct layers, red, green, and blue [40]. The computer saves the visible data in the TIFF format, preserving the dimensionality provided by the filter. When the visible image is read in software, it is recognized as a three dimensional matrix with the third dimension being the color layer. It is simple then to extract the red layer for use as the *Red* data in the NDVI calculation. A MATLAB script was written that when given a visible and NIR image calculates the NDVI pixel-by-pixel, and outputs an "NDVI image." An algorithm for forming an NDVI image, as well as sample output from MATLAB are shown below, in Algorithm 1 and Figure 4.3 respectively.

---

**Algorithm 1** Form NDVI Image
 

---

1.  $VIS \leftarrow \text{read}$  'visible image'
  2.  $NIR \leftarrow \text{read}$  'near-infrared image'
  3. Set *Red* equal to the first layer in *VIS*
  4. Form *NDVI* image by Equation 5.1  $\{(NIR - Red) ./ (NIR + Red)\}$
  5. **return**  $\rightarrow NDVI$  {greyscale image}
- 




Figure 4.3: NDVI image formed from Figure 4.2

It is pertinent now to discuss the bit depth used to present the NDVI image and its relation to the NDVI scale. The NDVI image is displayed in 8 bit grayscale, where a pixel can have a value ranging from 0 – 255, with 0 being pure black and 255 being pure white. The value of each pixel has been determined by the NDVI calculation described above, but by examining the pseudo-code used to form the NDVI image the output values can be observed to exist from 0 to 1. However, in representing this data as an 8 bit image the calculated values have been mapped to the 0 – 255 grayscale. In this way, the image can be properly displayed in software and remains human readable. For the remainder of this work, when the value of a pixel is discussed it will be referred to as the “pixel intensity”, and when it is desired to draw conclusions from NDVI images pixel intensity will be related back to the usual NDVI scale.

It becomes apparent that with the forming of an NDVI image, there is a wealth of information a human can discern immediately upon viewing the image, as opposed to a data vector provided by a spectrometer. The task of this work then is to interpret that information beyond what is first apparent in the image. It is common practice in the machine vision field to begin image analysis with a segmentation step. The goal of segmentation is to reduce or remove noise, and isolate the features of interest. In a way, forming the NDVI image has acted as an initial segmentation; the vegetation of interest appears bright and stands out, while some background elements that are not vegetation appear black. Unfortunately, NDVI is not enough and significant soil noise, as well as human elements (shoes specifically), still appear in the ground NDVI image. A simple threshold could be performed, eliminating features below a selected NDVI value. When this approach is used significant amounts of vegetation are lost, and considerable soil noise still appears. Therefore, a second segmentation stage was applied to remove the majority of soil noise, but retain as much vegetation as possible.

The second segmentation stage makes use of the Green NDVI, commonly abbreviated as GNDVI, discussed previously. Recall the general equation for GNDVI and that it relies only on data from the visible spectrum [41].

$$GNDVI = \frac{Green-Red}{Green+Red} \quad (4.2)$$

The *Green* data of the visible image is extracted in the same manner in which the *Red* data was, and a GNDVI image formed with the appropriate pixel-by-pixel operations. Upon inspection of

the GNDVI image it is apparent that much of the soil noise present in the NDVI image is missing from the GNDVI. This effect is used to create a binary mask, from the GNDVI image, to reject background noise. The mask is a matrix of the same dimensions of the image, where entries only have one of two values, 0 or 1. The value of each matrix entry is determined by a threshold set by the user. In the case of the ground images the threshold was set to twice the mean pixel intensity of the entire GNDVI image. This threshold was set experimentally and gave a suitable balance of noise rejection to vegetation retention. After the mask has been created it is applied to the NDVI image by matrix multiplication and a final segmented image is formed. The effect of the mask is that every pixel of the NDVI image is evaluated, and if the corresponding entry in the mask has a 0 value then the pixel in the final image is given a value of 0 as well, else the pixel in the final image retains its value from the NDVI image. The GNDVI helps to better recognize vegetation from noise, and that knowledge is applied to the NDVI image resulting in a final image which contains low noise and the majority of the vegetation present in the scene. Pseudo-code used for the entire segmentation process, as well as images from each step, are provided below, in Algorithm 2 and Figure 4.4 respectively.

---

**Algorithm 2** Two Stage Segmentation
 

---

1.  $VIS \leftarrow$  **read** ‘visible image’
  2.  $NIR \leftarrow$  **read** ‘near-infrared image’
  3. Set *Red* equal to first layer in *VIS*
  4. Set *Green* equal to second layer in *VIS*
  5. Form *NDVI* image by Equation 5.1  $\{(NIR - Red) ./ (NIR + Red)\}$
  6. Form *GNDVI* image by Equation 5.2  $\{(Green - Red) ./ (Green + Red)\}$
  7. Set *GreenLevel* threshold
  8. Create binary *mask* by  $GNDVI > GreenLevel$
  9. Form *SegmentedNDVI* image by applying mask to *NDVI*  $\{NDVI .* mask\}$
  10. **return**  $\rightarrow$  *SegmentedNDVI*
- 




Figure 4.4: Two Stage Segmentation: a) initial NDVI image, b) GNDVI image, c) mask, d) segmented NDVI image

The final image, referred to as the segmented NDVI image, contains data about the corn captured in the scene on the NDVI scale and is now suitable for analysis in regards to the nitrogen rates. To compare corn stalks from different subplots, a weighted mean pixel intensity is calculated over the segmented NDVI image. The equation used to calculate the weighted mean pixel intensity is as follows:

$$MPI = \frac{\sum_{i=1}^N p_i}{N}, \quad \forall p \neq 0 \quad (4.3)$$

Where  $p_i$  is the intensity value of a pixel, and  $N$  is the number of non-zero-valued pixels. The mean is weighted in the sense that it is the mean only calculated over the area of the image which is not black, i.e. where the pixel intensity is not equal to zero. It was chosen to weight the mean this way because there were different amounts of vegetation in each image. If a very bright corn stalk took up a small percentage of the image area, then the true mean pixel intensity of that image is still low, even though the vegetation in the image is bright. By not including the number of zero valued pixels as part of the mean calculation, the amount of vegetation in an image can vary without significantly skewing the results of the analysis. The mean pixel intensity (MPI) was chosen as the test variable because it is a simple and fast calculation that could be applied over a large set of test images. The MPI is also synonymous with the average NDVI of the vegetation in the scene, and will later allow conclusions to be drawn about how fertilizer rate affects NDVI.

### 4.1.3 Results

One set of visible and NIR images were selected per nitrogen rate for the P1184AM hybrid. The two stage segmentation described in the previous section was applied to each set of images and the weighted mean pixel intensity was calculated from the segmented NDVI images. Recall, the goal of the ground image study was to prove that the multispectral system can distinguish well fertilized corn from less fertilized by means of NDVI images. The results of the ground image study are shown below in Table 4.1.

Table 4.1: Mean Pixel Intensity and NDVI for one sample ground image test

<b>N Rate (lbs/ac)</b>	50	100	150	200
<b>MPI</b>	97.7	105	107	108
<b>NDVI</b>	0.382	0.410	0.418	0.422



The one sample test results show that there is indeed a difference in the mean pixel intensity, and thereby the NDVI, of well fertilized corn from less fertilized corn. The question now is if this trend is true over the entire test field. Although significant conclusions cannot be drawn from the one sample ground image test, confidence has been established in the potential of the system to use NDVI images as a tool to detect fertilizer rate and nutrient related stress of the crop. After the success of the ground image test, the fully integrated UAS is ready to fly over the entire test field. The flight images will be analyzed to answer the question of whether or not more nitrogen fertilizer always leads to higher NDVI and a healthier corn crop.

## 4.2 Flight Images

Now that the usefulness of the multispectral system had been proven with a ground image study, the complete unmanned crop monitoring mission was ready to be flown. Please see the previous chapter for details on the integration of the multispectral payload, as well as the unmanned vehicle used for the mission. The goal of the flight mission was to apply the multispectral system over the entire test field from an aerial perspective to detect the different fertilizer rates applied to the crop. Statistical inference will be used to prove what nitrogen rates can be reliably distinguished by the system. If the system can distinguish well fertilized subplots from the less fertilized, then it will have proven one way in which it can monitor crop status. The goal is to show a correlation between fertilizer rate and NDVI, thereby providing a direct comparison to the optimum nitrogen uptake estimates provided by Dr. Thomason. A secondary goal of the flight mission was to explore how sensitive the system was to changes in camera exposure. It is common practice in machine vision that the exposure of any system be specifically calibrated before use, as changes in exposure can lead to radically different results. Although a thorough test of camera exposure was not conducted, the hope is that the flight test results will be able to shed some light on how the system responds to different exposures.

### 4.2.1 Image Acquisition

The crop monitoring mission was flown on July 16<sup>th</sup>, 2013 in partly cloudy conditions. The corn was now 10 weeks old and approaching the tasseling stage, with the tallest stalks standing over 6 feet high. Three separate passes over the field were taken with camera exposures of 1000  $\mu$ s, 750  $\mu$ s, and 300  $\mu$ s respectively. The camera was triggered at 10 Hz and roughly 500 images in both the visible and NIR spectrum were captured for each pass. This provided more than enough data to analyze the effectiveness of the multispectral system. Altitude for the flight was set to 20 meters above ground level and the forward velocity to 1 m/s. With the camera's field of view fixed to 46°, almost 56 feet of ground was covered in each image and resulted in a resolution of 0.65 inches per pixel. Figure 4.5 shows a montage of a visible and NIR flight image. Right away, one advantage of the UAS is apparent in that it is able to fly much closer to the crop and provide significantly higher resolution than manned aircraft. For example, a study using manned aircraft imaging in south-eastern Australia achieved a max resolution of just 20 inches per pixel [42]. This increased resolution allows the system to potentially look at individual corn stalks and describe their status.




Figure 4.5: Visible (left) and NIR (right) image gathered during flight on 7/16.

After all the images were gathered they once again needed to be referenced to which subplots were visible in the image. Unlike the ground images, each flight image covered an entire range and subsequently two subplots, one of each hybrid. It will be important to distinguish between hybrids with the aerial images because each hybrid might react differently to the nitrogen rates. However, the referencing was easier with the flight images, as the aerial perspective allowed a

sanity check on which subplots were being captured in each image. Due to the high frequency in which the images were taken, there was an immense amount of overlap between each subsequent image. While this would be a boon for creating a mosaic of the field, it instead provides an abundance of repeated data for any statistical analysis. Therefore, one image of each range, from each pass, was selected for further analysis. The images selected had each range well centered in the image with a minimal amount of yaw, or rotation from the primary axis of the corn rows.

#### **4.2.2 Image Processing**

The method selected to analyze the flight images is an overall analysis of variance (ANOVA) followed by a Fisher's Least Significant Difference (LSD) test. Both ANOVA and LSD are popular statistical methods used to compare means of independent treatments and are well studied for use in agricultural sciences [43]. The overall test is used to first see if there is a detectable difference between any of the selected treatments, and the LSD test used after to determine if each treatment is detectable from every other. If the overall ANOVA is unable to show any difference in the treatment means then the LSD test will not be carried out. In this way, ANOVA acts as initial screening of the data and "protects" the LSD test from a high false positive rate. This method is commonly referred to as a Fisher's Protected LSD [44]. The protected LSD is more likely to identify differences in treatments than other statistical comparisons, but at the cost of a higher rate of false positives. The higher false positive rate is an acceptable risk in order to prove that the multispectral system can differentiate between the different nitrogen rates. Three protected LSD tests are presented here in this section: first to test the effects of camera exposure on the overall system performance, second to test for detection of the N rates in the P1184AM hybrid, and lastly to test for N rate detection in the P1745HR hybrid.

The statistical inference begins with establishing the null and alternative hypotheses. The null hypothesis, which needs to be disproved, states that there will be no observed effect for the test, and the alternative hypothesis states what effect will be observed. The overall ANOVA test is used to provide evidence for disproving the null hypothesis, confirming that some effect does exist for causing differences in treatment means, and the LSD test is used to explore those differences further. For the flight image analysis the hypotheses can be succinctly stated as follows:

$$\begin{aligned}
 H_0 &= \mu_A = \mu_B = \mu_C = \mu_D \\
 H_a &= \text{at least one mean differs}
 \end{aligned}
 \tag{4.4}$$

where  $H_0$  and  $H_a$  are the null and alternative hypotheses respectively, and  $\mu_{A-D}$  are the population means for the treatments under study. An initial test was done to look at the effect of exposure on the system, in which case the statistical treatments are the three test exposures of 1000  $\mu\text{s}$ , 750  $\mu\text{s}$ , and 300  $\mu\text{s}$  respectively, and there is no fourth treatment. When testing for the effect of exposure corn captured in the flight images is divided into three populations depending on what exposure was used for a given image. For the main body of analysis, the statistical treatments are the four nitrogen rates of 50 lb/ac, 100 lb/ac, 150 lb/ac, and 200 lb/ac respectively. It follows then, that the corn is divided into four populations dependent upon which nitrogen rate was applied to a given subplot.

A confidence level also needed to be established before any inference testing took place. The confidence level chosen was 95%, and the resulting alpha value is 0.05. Alpha is the maximum probability that a null hypothesis is rejected when it is in fact true. The alpha value also controls the rate of false positives, which in the case for this project is maximally five for every one hundred observations, or simply put 5%. It is important to note that alpha is the maximum probability for an observation being a false positive and is not a definite probability. As stated earlier, Fisher's LSD has a higher rate of false positives than other multiple comparison methods, and it has been decided that the maximum rate allowed for the project will be 5%. The effect of the chosen confidence level is that up to 5% of all the corn may be detected as part of the incorrect nitrogen rate, and that the system will observe the test results 95% of the time.

The key to the LSD statistical inference method is that the data must be gathered by a simple random sample (SRS). Therefore, a method had to be developed in which the images could be sampled randomly. First, a sample size had to be selected. Would the entire subplot be used as a sample, would each pixel be used as a sample, or what about a grouping of pixels? Each sample size could fit well into a variety of tests, but it is most sensible that a body of pixels be sampled for the flight image analysis. The chosen sample size was 40 pixels wide by 50 pixels tall, or equivalently 26" x 33" on the ground. Figure 4.6 shows this sample size in respect to a typical

flight image. This sample size was the most reasonable for covering a single corn stalk that does not overlap significantly with its neighbors, and in the worst case, only covers two corn stalks that are growing closely together. In this way, it is possible to relate the sample size to the amount of canopy coverage the average corn stalks possesses.



Figure 4.6: A single sample (pink) has been selected from a visible flight image.

As with the ground images, the vegetation in the flight images needs to be segmented before sampling. Therefore, the first step of the sampling routine is to calculate the NDVI of the selected images. In the case of the flight images, a single stage of NDVI segmentation was sufficient to separate the corn from any background noise. Recall that a second stage of segmentation was required with the ground images due to the presence of soil noise and human artifacts in the scene. When the flight images were gathered the corn had grown to a point where many leaves of a given corn stalk reached over the gaps between its neighboring stalks. This was true for stalks in the same row as well as stalks in adjacent rows. The result was minimal soil noise present in the flight images, and after the formation of a NDVI image the majority of all soil appeared black effectively segmenting the corn from the noise. The only artifacts present in the flight images were the shadow of the UAS as well as the irrigation lines. However, each of

these artifacts is easily segmented as part of the background by NDVI, as neither has the spectral response to be characterized as vegetation. Figure 4.7 shows a flight image highlighting a shadow, and the resulting NDVI image.




Figure 4.7: A visible flight image (left) and the resulting NDVI image (right). Notice how the shadow (red) in the visible image is no longer present in the NDVI image.

With the vegetation segmented from the background noise, the next step to sampling the data for statistical analysis was to identify the N rate and hybrid variety of the corn in each image. Recall that each flight image used for analysis captures a single range, and that each range is divided into two subplots, one of each hybrid. A user is prompted by the software running the sampling routine to identify the hybrid variety of the subplot in the left half of the image, followed by the N rate of that subplot, and finally to draw a bounding box around the subplot. This process is completed for each the flight images selected for analysis. The bounding box is constrained to a regular rectangle for ease of use in forming a grid of samples. The drawback with a rectangular bounding box is that if the helicopter has yawed relative to the primary axis of the corn rows the subplots captured in the image will not line up well with the bounding box. In this case, the user is faced with the decision of creating a bounding box that either contains areas of the background, or misses portions of the subplot. Although this case was rare with the flight images used for sampling, preference was given to a bounding box that covered the majority of the subplot while minimizing the inclusion of the background. Each image covers more than enough area for sampling, and therefore it was deemed more important to reject as much background from the samples as possible then attempt to cover the entire subplot. In this way, the data

calculated from the samples will be more representative of the entire subplot than that of the background. Figure 4.8 shows a NDVI image with a user drawn bounding box.




Figure 4.8: A bounding box (blue) drawn by the user during the sampling routine. The bounding box is drawn to identify one of the two subplots in the image.

After the bounding box has been drawn by the user, the sampling routine creates a regular grid within the bounding box. The cell size of the grid is that of the selected sample size, and only whole cells are used with any remainder of the bounding box ignored. Each cell is now a potential sample to be used for statistical inference, and to satisfy the SRS requirement random cells from the grid are selected as samples each time the routine is run. This is all done by software once the user has drawn the bounding box, and an example of a sampled image can be seen in Figure 4.9. For the test on exposures 12 samples were taken from each subplot, and for the test on nitrogen rates 15 samples were taken per subplot. This method for sampling the flight images satisfies the randomness requirement. Even though relatively few flight images were selected for analysis, each image yields over 24 data points, which is more than enough for significant statistical inferences to be made.




Figure 4.9: After bounding box selection random samples are chosen from each subplot. The first subplot and the corresponding samples are outlined in green, and the second subplot and its samples in red.

During the sampling routine, as each sample is selected from the image, the mean pixel intensity of the sample is calculated and recorded. It was once again decided to calculate MPI by Equation 5.3. Recall that Equation 5.3 is not a straightforward mean calculation, as it does not include pixels with a zero intensity into the number of observations used to divide the sum pixel intensity by. Succinctly put, the mean pixel intensity is only calculated over the foreground and neglects the background (i.e. black) pixels. For the ground images it made sense to neglect zero valued pixels because it allowed the size of the corn stalks to vary between the images while preserving the ability to extract the average NDVI of the stalk. In the case of the flight images, there are additional reasons to neglect black pixels in the mean. Examine the two samples as shown in Figure 4.10. Both samples are of the same size and are from the same subplot, however the sample on the left has significantly more black pixels than the sample of the right. One could make the argument that the number of black pixels in the sample relates to decreased canopy coverage and is a symptom of a nutrient deficient plant. However, background pixels could be introduced into a sample by various means such as: roll and pitch of the helicopter, uneven row



spacing, and the slope of the test field. These effects are difficult to quantify without further testing, and therefore background pixels will continue to be neglected in the MPI calculation. Also, due to randomness of the grid based sampling method, a sample could be selected on edge of the subplot, or in the worst case, in between rows. By disregarding the black pixels potentially introduced by random samples, the quick grid based approach can be used without concern for skewing the data.




Figure 4.10: Two samples taken from the same subplot in the same image. Sample a) has significantly more black pixels than sample b) and without proper consideration sample a) could appear to be from a less fertilized subplot.

To illustrate the effectiveness of Equation 5.3, the mean pixel intensity of the two samples shown in Figure 4.10 is calculated by two methods: first by a simple mean calculation, and secondly by the weighted mean calculation. The results of the calculations, as well as percent differences, are shown in Table 4.2. The percent difference from sample a) to sample b) is almost double when using the simple mean method. With this one sample example it might appear that that a different N rate was used on each of the samples. If the simple method were used over the entirety of the flight images expect greater variability in the MPI of the samples could be expected, as well as a general skew of lower pixel intensity. Instead, by using the weighted mean method shown in Equation 5.3, the variability in MPI of the samples is decreased and the ability to detect differences in nutrient status caused by N rate is increased.

Table 4.2: Mean Pixel Intensity calculated by two methods

<b>Sample</b>	a)	b)	Percent Difference (%)
<b>Simple Mean</b>	140	152	7.14
<b>Weighted Mean</b>	153	159	3.92

The entire sampling routine from a set of visible and NIR images to MPI is described below in Algorithm 3.

---

**Algorithm 3** Sample Flight Image

---

1. Set *SampleWidth* and *SampleHeight*
  2. Select *N* number of flight images
  3. Select *S* number of samples per image
  4. **for** *i* = 1 to *N*
  5.     *VIS* ← **read** ‘visible image’
  6.     *NIR* ← **read** ‘near-infrared image’
  7.     Form *NDVI* image from Algorithm 1.
  8.     Request user input for Hybrid type and Nitrogen rate
  9.     Request user draw *BoundingBox*
  10.    Set *ColumnNumber* equal to *BoundingBox(WIDTH)/SampleWidth*
  11.    Set *RowNumber* equal to *BoundingBox(HEIGHT)/SampleHeight*
  12.    Create regular *Grid* from *ColumnNumber*, *RowNumber*, *SampleWidth*, *SampleHeight*
  13.    Select *S* random *Grid* locations
  14.    **for** *j* = 1 to *S*
  15.       Form *Sample* by cropping *Grid(j)* from *NDVI*
  16.       Calculate MPI of *Sample* by Equation 5.3
  17.    **end for**
  18. **end for**
- 

As implemented for this work, lines 8 – 17 of Algorithm 3 were executed twice per flight image due to two subplots being captured in every picture.

### 4.2.3 Exposure Test Results

Before testing the flight images for nitrogen rate detectability, it was desired to test the effects of camera exposure on the mean pixel intensity of the field. As mentioned previously, a common drawback to visual sensors is that they must be calibrated before each test, and this includes setting an exposure to adjust the overall brightness of the collected images. It is preferable that payloads for UAS not require frequent calibrations so that they may be ready to fly a mission whenever desired. The ideal visual sensor would be insensitive to changes in exposure and produce comparable results after every mission; however, this is almost never the case with visual sensors. Therefore, it was pertinent to explore how sensitive the multispectral system is to changes in exposure. The first protected LSD test compares the mean pixel intensity of the entire corn crop with each exposure used to capture the field, which can then be related to the average NDVI of the field. The LSD test will provide answers to how camera exposure affects the mean

pixel intensity of the NDVI images and if samples from each exposure can be combined for the nitrogen rate detectability tests.

The first step of the protected LSD test is an overall ANOVA. The overall ANOVA is done to disprove the null hypothesis before looking at the differences between each treatment mean. In the case of this exposure test, the treatments are the three exposures used in gathering the flight images, and the treatment means are calculated as the MPI of all the samples taken with each exposure. Recall, the null hypothesis for this work states that each treatment mean is equal, i.e. there are no differences between any of the treatments. In terms of the exposure test the null hypothesis says that a change in exposure will not create a significant difference in the mean pixel intensity of NDVI images, and that the average NDVI of the entire field for each exposure setting should remain the same. In order to disprove the null hypothesis the p-value of the test must be less than the predetermined alpha value. The p-value can be found in the far right column of the ANOVA table shown below in Table 4.3. For an explanation of each term in the ANOVA table please see [44]. Recall the confidence level chosen for every test was 95% and the resulting alpha level is 0.05. Therefore, with a p-value of less than 0.0001 the overall ANOVA has disproved the null hypothesis and proves that at least one treatment mean differs from the others. For the exposure test this is an unfortunate result as it means the multispectral system is sensitive to changes in exposure, and consequently, the flight images gathered with at least one of the exposure settings is not comparable to the others.

Table 4.3: ANOVA of exposure test results.

Source	Degrees of Freedom	Sum of Squares	Mean Square	F Ratio	p-value
Exposure	2	37501.42	18750.7	69.51	< 0.0001
Error	789	212836.38	269.8		
Total	791	250337.80			

Now that exposure has been proven to cause a significant difference in a sample's MPI for at least one case, the Fisher's Least Significant Difference test is used to examine how the mean of each exposure treatment differs from the others in a pairwise sense. Two tools are used for describing the results of the LSD test: a Connected Letters Report, and confidence intervals. A Connected Letters Report assigns a letter (e.g. A, B, C, etc.) for each grouping of treatments that

have not been proven to be significantly different. Any treatment assigned only a single letter is believed to be significantly different than the other treatments and can be reliably distinguished. Confidence intervals are used to describe a belief in where the numeric difference between two population means lies. Each interval has an upper and lower bound determined by the confidence level chosen for the test and the standard error calculated for the data. If the mean of the treatment listed first is believed to be greater than the second the entire confidence interval will be positive (i.e.  $\mu_i - \mu_j > 0$ ), and the inverse can be said if the first treatment mean is believed to be less than the second. If the difference between two treatment means is not found to be significantly different then the confidence interval for those treatments will contain zero. This demonstrates a belief that the true difference between the population means is equal to zero (i.e.  $\mu_i - \mu_j = 0$ ). Also listed for each interval is a p-value, and similar to the ANOVA test, if the p-value is less than the selected alpha level the treatments cause some effect to occur and the two population means are believed to be significantly different. For this work, confidence intervals are shown in table format and have been calculated with a 95% confidence level.

Table 4.4 displays the Connected Letters report for the LSD test on camera exposures. The report shows that the top two exposures (1000  $\mu\text{s}$  and 750  $\mu\text{s}$ ) are connected and not significantly different. However, the darkest exposure is not connected to the others and the mean for that exposure setting is significantly lower than the other two. Beyond observing the means in Table 4.4 the confidence intervals in Table 4.5 can be inspected. As expected, the interval for the difference between the first and second exposures includes zero, and therefore the population means of each pass are believed to essentially be equal. This means that changing the exposure from 1000  $\mu\text{s}$  to 750  $\mu\text{s}$  did not have a significant effect on the mean pixel intensity of the NDVI images, and that each pass has detected the same average NDVI for the entire field. The other two intervals have almost the exact same upper and lower bounds and show that when the exposure was changed to 350  $\mu\text{s}$  the mean pixel intensity of the NDVI images decreases significantly. The average NDVI for the entire field should not change significantly from pass to pass, but the results of the third pass show a decrease in field NDVI by roughly 0.057 from the first two passes. It can be concluded that the decreased brightness in the visible and NIR images gathered with the lowest exposure formed NDVI images of lower mean pixel intensity, and skewed the average NDVI of the field lower than its true value. Due to the difference in MPI

caused by exposure, the images taken with the 350  $\mu\text{s}$  exposure will be neglected in the analysis of nitrogen rates. The decreased brightness of those images could skew the results to report lower means for each N rate and potentially obscure differences between the rates.

Table 4.4: Connected Letters Report for exposure test.

<b>Exposure</b>	<b>Connection</b>	<b>MPI</b>	<b>Average NDVI</b>
1000 $\mu\text{s}$	A	139	0.543
750 $\mu\text{s}$	A	139	0.543
350 $\mu\text{s}$	B	124	0.484

Table 4.5: Confidence Intervals for exposure test.

<b>Treatment</b>	<b>Less Treatment</b>	<b>Lower Bound</b>	<b>Upper Bound</b>	<b>p-value</b>
1000 $\mu\text{s}$	750 $\mu\text{s}$	-2.52	3.10	0.840
1000 $\mu\text{s}$	350 $\mu\text{s}$	11.9	17.5	< 0.0001
750 $\mu\text{s}$	350 $\mu\text{s}$	11.6	17.3	< 0.0001

As previously mentioned, the exposure test was completed only to see the effect of changing exposure on the pixel intensity of NDVI images. It is not possible at this time to state which exposure setting provides more accuracy on the NDVI scale and could be considered “correct.” To do so a calibration of the multispectral camera would need to be carried out with sources of known reflectance in the red and NIR bands. In this work, the specific NDVI value of a plant is not of interest, but instead the difference in the NDVI of plants at different nutrient statuses. Although it is not possible to state what exposure setting produces the most accurate NDVI images, the main body analysis relies only on differences in pixel intensity to distinguish well fertilized corn from the less fertilized. An argument could be made, that without a better understanding of the system’s sensitivity to exposure, and a specific calibration, the results of the various tests are not repeatable. However, because the analysis relies on differences in NDVI, and not the NDVI value itself, it is the opinion of the author that a calibration is not necessary as those differences of interest should exist as long as the resulting NDVI images are not saturated.

#### 4.2.4 N Rate Detectability Test Results

After the conclusion of the exposure test, the flight images were resampled for the nitrogen rate detectability tests. Thirty (30) samples were taken from each image, 15 per subplot, for a total of 90 from each N rate. Two detectability tests were completed, one for each of the two corn hybrids present in the field. The tests followed the same protected LSD method as that used for the exposure test. The treatments for these tests are the four nitrogen rates applied to the field, and once again, the treatment means are calculated as the MPI of all samples collected from the flight images. The goal of each detectability test is to determine if NDVI images can consistently distinguish the well fertilized corn from the less fertilized. The ground image study showed a positive correlation between average NDVI and fertilizer rate. But the question that remained is, “will greater applications of fertilizer always lead to higher NDVI, or does there come a point where increasing the amount of fertilizer fails to result in an increase of NDVI?” Finally, each test should also form a baseline estimate for the average NDVI of healthy corn versus nutrient deficient corn.

The first hybrid studied was the P1184AM, and following with the protected LSD method an overall ANOVA was carried out. Table 4.6 shows the results of the overall ANOVA for the P1184AM hybrid. The p-value for the ANOVA is reported as less than 0.0001 and is considerably smaller than the selected alpha level of 0.05. A p-value that small states there is overwhelming evidence to reject the null hypothesis, and confirms that at least one treatment mean differs significantly from the others. This is the expected result, as it has been previously observed that the MPI of NDVI images changes with respect to the applied fertilizer rate. The overall ANOVA has disproved the null hypothesis, so the next step is to move on to the least significant difference test.

Table 4.6: ANOVA of P1184AM nitrogen rate detectability results.

Source	Degrees of Freedom	Sum of Squares	Mean Square	F Ratio	p-value
N Rate	3	17701.63	5900.54	17.85	< 0.0001
Error	356	117665.85	330.52		
Total	359	1353670.48			

A Connected Letters Report for the P1184AM hybrid presented in Table 4.7 shows that only one nitrogen rate can be reliably distinguished from the rest. The top two N rates (200 and 150 lb/ac) have been grouped together as have the middle two rates (150 and 100 lb/ac). Only the lowest N rate of 50 lb/ac has been deemed significantly different from all the other rates. Table 5.8 shows that all but two of the confidence intervals for the first hybrid contain zero. The p-values for the unconnected treatments are all much less than the alpha level and indicate a considerable confidence that the difference between the treatment means is greater than zero. Conversely, the p-value for the interval of the top two nitrogen rates is more than five times the alpha level and indicates a strong belief that there is no difference in the population means. Most interesting is the confidence interval for the 150 lb/ac treatment to the 100 lb/ac. The p-value for the above mentioned interval is 0.0630, very close to the alpha level of 0.05, and the interval itself just barely includes zero. If the data were to be resampled, and the test run again, there is a high probability that the 150 and 100 lb/ac treatments could be found significantly different. Also, if the test were run with a smaller confidence level, a common choice would be 90%, it is almost certain that the middle two treatments would be deemed significantly different. As it stands, only the P1184AM corn sidedressed with 50 lb/ac of nitrogen is independent of the rest and could consistently be detected with 95% confidence by the multispectral system.

Table 4.7: Connected Letters Report for P1184AM N rate detectability test.

<b>N Rate</b>	<b>Connection</b>	<b>MPI</b>	<b>Average NDVI</b>
200 lb/ac	A	145	0.566
150 lb/ac	A B	142	0.554
100 lb/ac	B	137	0.535
50 lb/ac	C	127	0.496

Table 4.8: Confidence Intervals for P1184AM N rate detectability test.

<b>Treatment</b>	<b>Less Treatment</b>	<b>Lower Bound</b>	<b>Upper Bound</b>	<b>p-value</b>
200 lb/ac	150 lb/ac	-2.32	8.34	0.268
200 lb/ac	100 lb/ac	2.73	13.4	0.0031
200 lb/ac	50 lb/ac	13.1	23.8	< 0.0001
150 lb/ac	100 lb/ac	-0.276	10.4	0.0630
150 lb/ac	50 lb/ac	10.1	20.8	< 0.0001
100 lb/ac	50 lb/ac	5.06	15.7	0.0001

The second corn hybrid, P1745HR, was studied in the same manner as the first. The overall ANVOA table is shown below in Table 4.9. The result of the overall ANVOA is similar to that of the first detectability test; a p-value much less than the alpha level, and overwhelming evidence that one treatment mean differs significantly from the others. Due to the results of the P1184AM hybrid, the LSD test can be expected to, at the very least, show strong confidence that corn sidedressed with 200 lb/ac of nitrogen has a significantly different MPI than corn sidedressed with 50 lb/ac.

Table 4.9: ANOVA of P1745HR nitrogen rate detectability results.

Source	Degrees of Freedom	Sum of Squares	Mean Square	F Ratio	p-value
N Rate	3	11093.08	3697.69	15.37	< 0.0001
Error	356	85645.51	240.58		
Total	359	96738.60			

The Connected Letters Report for the P1745HR hybrid offers more promising results than the first hybrid. The report, presented in Table 4.10, shows that two of the four N rates can be reliably distinguished from the others. Once again, the two highest N rates remain connected and cannot be considered significantly different. The confidence intervals displayed in Table 4.11 show that the interval for the 200 and 150 lb/ac rates not only contains zero, but has a p-value much greater than the 0.05 alpha level used for the test. The 200 – 150 lb/ac interval conveys an overwhelming belief that the two treatments have equal population means. All of the other confidence intervals show the converse relationship, that there is a strong belief the population means are significantly different. Unlike the P1184AM hybrid, the 150 lb/ac and 100 lb/ac N rates produced significantly different means without any need for resampling or lowering the confidence level. The P1745HR corn hybrid has shown a similar response to fertilizer rate as the P1184AM hybrid, but with an increased ability to distinguish subplots sidedressed with 100 lb/ac of nitrogen.

Table 4.10: Connected Letters Report for P1745HR N rate detectability test.

Treatment	Connection	MPI	Average NDVI
200 lb/ac	A	154	0.602
150 lb/ac	A	154	0.602
100 lb/ac	B	147	0.574
50 lb/ac	C	141	0.551



Table 4.11: Confidence Intervals for P1745HR N rate detectability test.

<b>Treatment</b>	<b>Less Treatment</b>	<b>Lower Bound</b>	<b>Upper Bound</b>	<b>p-value</b>
200 lb/ac	150 lb/ac	-4.01	5.08	0.817
200 lb/ac	100 lb/ac	2.87	12.0	0.0015
200 lb/ac	50 lb/ac	9.00	18.1	< 0.0001
150 lb/ac	100 lb/ac	2.34	11.4	0.0031
150 lb/ac	50 lb/ac	8.46	17.6	< 0.0001
100 lb/ac	50 lb/ac	1.58	10.67	0.0084

Least significant difference testing has proved that the multispectral system is able to distinguish well fertilized corn from less fertilized in the most extreme case for both hybrids. Both detectability tests gave overwhelming evidence that fertilizer rate has an effect on crop NDVI, and that corn sidedressed with nitrogen rates above 150 lb/ac consistently have higher average NDVI than corn sidedressed with only 50 lb/ac of nitrogen. The average NDVI for the 50 lb/ac N rate for each hybrid can then be used to form a baseline for segmenting corn based on its nutrient level. If there is confidence in the calculated NDVI for each sample, the multispectral system can make a binary decision as to whether or not a crop is nutrient deficient based on the deviation from the baseline. For example, a sample of P1745HR corn with an average NDVI above 0.56 can be labeled as healthy relative to a sample with average NDVI below 0.56. Increased confidence in the NDVI calculations will come from a proper calibration of the multispectral camera. Following that calibration, the multispectral system could then be used in the future for field-wide nutrient deficiency monitoring.

#### **4.2.5 Grain Yield Comparisons**

The final result explored is a comparison between the multispectral system results and the field results gathered by Dr. Thomason. Specifically, it was desired to relate the trends of average NDVI calculated from the flight images and grain yield measured after harvest to the nitrogen rates. The test field was harvested in mid-October and grain yield for each subplot measured, and then the subplot yields summed depending on their hybrid type and nitrogen rate. Figure 4.11 shows the measured grain yields plotted against the applied nitrogen rates. Also displayed in Figure 4.11 are lines of best fit calculated in the typical least squares sense. The first thing to notice is the yield plateau reached by the P1174HR hybrid. Corn of that hybrid did not see an increase in grain yield when increasing the fertilizer rate from 150 lb/ac to 200 lb/ac. The plateau

and the quadratic nature of the yield response suggest that optimum N uptake was reached by the corn, and exists somewhere between 150 and 200 lb/ac. On the other hand, no plateau was reached for the P1184AM hybrid and there is only a semi-linear relationship is noticeable between N rate and grain yield. As with any field study, many factors affect yield and it is perhaps the case that an unaccounted for variable has interfered with the fertilizer rate study of the P1184AM corn. Regardless, looking only at the data as presented it is most likely the case that optimum N uptake either occurred at the maximum N rate or beyond the range of N rates used. The optimum N rate for the P1745HR hybrid can be quickly estimated by finding the peak of the corresponding trend line. This is done in the usual way by taking the first derivative of the quadratic trend line and solving for the zero crossing. After said method is complete, the result is an estimate of 161 lb/ac of nitrogen for optimum uptake. A similar estimation is not possible with the other hybrid due to its mediocre correlation with grain yield.




Figure 4.11: Grain yield versus applied Nitrogen rate

From Figure 4.11, a distinct difference in the nitrogen use efficiency of the two hybrids can also be observed. Nitrogen use efficiency is derived from the increase in yield in response to nitrogen application [45]. A corn hybrid with low nitrogen use efficiency will not see significant increases in yield as more nitrogen fertilizer is applied. For this study, the P1184AM hybrid has low nitrogen use efficiency, as demonstrated by an increase in grain yield of only 470 lb/ac even

when the applied N rate was quadrupled from the original rate of 50 lb/ac. In comparison, the P1745HR hybrid has a much greater nitrogen use efficiency than the AM hybrid. By simply doubling the baseline N rate of 50 lb/ac to 100 lb/ac, the HR hybrid saw an increase in grain yield of 820 lb/ac. It is safe to say then, that the P1745HR hybrid has a strong nitrogen use efficiency trait, and offers increased yield over the AM hybrid while requiring less nitrogen input.

The only question that remains is how the average NDVI calculations compare to the grain yield measurements. Figure 4.12 shows the grain yield measurements plotted against their respective average NDVI values. The plots show that NDVI has a strong correlation ( $r^2=0.82$ ) to grain yield for the P1745HR hybrid, but the correlation for the other hybrid is much less impressive ( $r^2=0.40$ ). As mentioned in the second chapter, NDVI has been used in the past to estimate yield with mixed results. The limiting factor in most yield estimations using NDVI was soil interference. However, the methodology used in this work to calculate NDVI is much more robust with respect to soil noise, so that is not a fair argument for poor correlation. Until more data is gathered and more studies complete, the mixed results presented in this work and seen previously by other researchers will have to suffice. In any event, it is possible that in the future the multispectral system can act as a useful tool for estimating grain yield with field-wide NDVI calculations.




Figure 4.12: Grain yield versus NDVI

# Chapter 5

## Late-Season Tobacco Study

The Unmanned Systems Lab collaborated with members of the Virginia Tech Agricultural Research and Extension Center (AREC) and Altria Inc. to complete a study of late-season tobacco with the multispectral unmanned aerial system. This chapter will focus primarily on the processing of images gathered during the single test flight, but will also outline the development of the study. The layout of this chapter will differ from the previous ones due in part to the way in which the tobacco study developed. Unlike the crop monitoring mission over the corn field, the tobacco study was seized “in the moment” as an opportunity by all parties to expand the use of the multispectral system and investigate its potential for tobacco monitoring. An overview of the test flight is given including: the obstacles overcome to conduct unmanned operations at a new location, the goal of the test flight, the location of the test flight, and the flight plan. The remainder of the chapter consists of the techniques and processes used to analyze the images gathered during the test flight. A method for determining the space between crop rows is detailed, as well as the specific use of NDVI for segmentation and stress estimation of the tobacco. Also presented is a secondary segmentation routine, which opens up a variety of new analysis options. This secondary segmentation is the final material covered in the chapter and describes a method to distinguish healthy green leaves from bleached leaves.

## 5.1 Flight Mission Overview

In May of 2013 serious consideration was given to the idea of a test flight of the multispectral system over a tobacco farm in South Hill, VA. Dr. David Reed from the AREC suggested this location as it was isolated from the surrounding area. Another benefit of flying over the suggested farm was that additional data, such as soil content analysis, could potentially be made available after harvest. Dr. Reed's suggestion was taken and the R. Hart Hudson tobacco farm selected as the test site for the study. The farm itself covers approximately 400 acres and is located along the north bank of the Roanoke River, 15 minutes south of South Hill. The range of soil types and topography of the farm allows for a wide variety of growing conditions and would give the project its choice of tobacco to study. Unfortunately, there existed an obstacle that would delay a test flight and restrict any flights to late in the growing season. Before any unmanned operations are conducted, permission from the FAA must be obtained in the form of a Certificate of Authorization (COA) for Unmanned Operations. Completing the COA was a straightforward task in of itself because the USL already has several COAs, and more specifically one for the project UAV, the RMAX. The real obstacle was that once completed and submitted to the FAA, a COA takes 60 days to be accepted. The timeline afforded to the COA process allowed only enough time for one test flight to be completed, and for only late-season tobacco to be imaged. The maturity of the tobacco will come into play in the next section about image analysis. In the end, the proposed COA was accepted by August and a flight test at the Hudson tobacco farm planned for September 14<sup>th</sup>.

Once a date had been selected for the flight over the tobacco, a mission plan had to be created in the same way it was for the corn monitoring mission. Dr. Reed identified a particular field at the Hudson farm that had yet to be harvested and was showing various signs of stress. It was then decided that the goal of the tobacco study would be estimating which block of tobacco rows were most stressed. The field selected by Dr. Reed was seven acres large; however, it was not desired to fly over the entire field. So, a flight plan was created that would span the entire width of the field (roughly 200 ft), but only cover the first 200 ft of its length. Figure 5.1 presented below shows the flight plan overlaid in Google Earth (© 2013 Google). Similar to the flight in July, the altitude was set to 20 meters AGL to achieve the maximum safe resolution captured by the multispectral system. Because the flight plan was to cover a significant portion of the field, the

flight speed was increased to 2 m/s between waypoints to decrease the flight plan time to just under 20 minutes. After the crew arrived at the farm, an appropriate position for the ground station and takeoff/landing zone was selected. Positioned by a nearby road and facing down the length of the field, the ground station had a complete view of the vehicle during the mission. The landing zone was located in front of the ground station and gave ample room for the helicopter to achieve altitude and be positioned before executing the mission flight plan. Overall, the test flight was a success with no damage sustained to personnel or equipment, and images of the tobacco field being gathered for post-processing.




Figure 5.1: Tobacco study mission plan

## 5.2 Image Analysis

The stated goal of the tobacco test flight is to identify the block of tobacco rows with the greatest estimated stress. Before estimating the stress of the crop, a method for determining the spacing between rows is presented. It was essential to have a solid estimate of the row spacing in order to identify the block of tobacco under the most stress. The row spacing estimate is then used to create boundaries around individual tobacco rows to allow for row by row stress estimation. NDVI is used to segment the tobacco from the background and as a point estimate for crop

stress, with stressed crops having lower NDVI. After an average NDVI has been established for each row of interest, a moving average is applied to identify the eight consecutive rows with the lowest NDVI. The result of the moving average is then presented as the estimate for the most stressed block of tobacco. The final image processing technique covered is a secondary segmentation method. Due to the maturity of the tobacco during the flight test, many of the tobacco leaves were “bleached,” and it was desired to develop a segmentation process that would allow for analysis of either just the green healthy leaves, or the bleached leaves. By inspecting the visible response of the bleached leaves, the secondary segmentation method was developed, and an example of its usefulness is presented to close the chapter.

### **5.2.1 Image Acquisition**

The tobacco test flight occurred on September 14<sup>th</sup>, 2013 in clear, sunny conditions. The tobacco field selected for the flight was well past maturity and showed visible signs of stress in the number of bleached leaves on each tobacco stalk. For the tobacco test flight the camera was only triggered at 1 Hz. This was acceptable as shown earlier by the excess of data provided when triggering at 10 Hz during the corn crop monitoring mission. The camera exposure was set to 400  $\mu$ s due to the ambient brightness caused by the sunny conditions. Figure 5.2 shows a sample image gathered during the flight and demonstrates that even though the exposure was low, the image is still bright and captures enough detail without washing out any of the vegetation. The flight plan described in the previous section was executed a single time, and was successful in capturing 600 images in both the visible and NIR spectrums. After the flight was complete, handheld GPS measurements were made of areas in the field that were evaluated by Dr. Reed to contain stressed tobacco. This was instrumental in limiting the data used for analysis. The area outlined by GPS measurements included the first 16 rows of tobacco and was approximately 50 feet down the length of the field. By correlating the handheld GPS measurements to the GPS log created by the system during flight, images could be selected that only looked at the tobacco of interest. The images matching closest to the GPS measurements were pulled from the data set and used in post-processing to identify the most stressed tobacco rows.




Figure 5.2: Sample tobacco study image.

### 5.2.2 Row Space Estimation

Essential to the tobacco study was estimating the stress of each individual row. To this end, a method was developed to analytically determine the average space between the rows. By knowing the row spacing, software can be written that segments an image into rows and calculates the mean pixel intensity of each row. This will allow for a row by row analysis and the identification of the most stressed block. Many machine vision techniques exist for identifying particular characteristics of an image, like intensity gradients or straight lines. However, a simpler method based on a fundamental of signal processing, the Discrete Fourier Transform (DFT), was selected. The DFT is used traditionally to transform discrete time domain data into the frequency domain [46]. There exists an analog when the transform is applied to spatial data, in that the resulting domain represents excitations that are occurring per unit distance, as opposed to per unit time [47]. When applied to an image, the DFT can be used across both of the spatial axes to investigate the frequency of significant changes in intensity. This lends itself nicely to determining the tobacco row spacing, because a row of tobacco will consist of an area of bright vegetation, followed by an area of darker ground. This pattern should then repeat for evenly planted rows. Taking a closer look at the sample image in Figure 5.2, it is obvious that although



the vegetation in a row does not have a consistent intensity, the most extreme change in pixel intensity occurs when moving from row to row. Therefore, a dominant frequency should appear in the DFT, and be analogous to the length of a row measured in pixels.

The first step in the DFT method is sample the image for data points. Each sample could be a single pixel, but is more likely to be a two-dimensional grouping of pixels, in which case the intensity of the group would be averaged to form a single data point. Samples that align on the spatial axis of interest are referred to as data vectors, and the entire collection of data vectors referred to as the data ensemble. To determine the row spacing, the DFT will only need to be conducted in the cross-row direction. Because the rows are orientated horizontally in the images, the axis of interest will be the vertical, or  $Y$ , axis. The vertical axis is where the change between rows occurs, and the DFT will provide an estimation of how many changes occur per pixel. DFT estimation is enhanced by performing multiple DFTs' and averaging their results together. In doing so, the noise present in each individual spectrum should average close to zero, and the dominant frequency should become more apparent. A reasonable number of averages should be selected, so as not to dramatically increase computation time.

The initial data ensemble used 10 averages to estimate the row spacing, and subsequently the ensemble would need 10 data vectors of even width. By inspecting Figure 5.2, it is apparent that the tobacco reaches all the way across the width of the image (1024 pixels). If the data ensemble is desired to cover the majority of the tobacco, a simple choice for data vector width is 100 pixels. The only choice that remains then is to decide the length of each sample. The Nyquist criterion states that for discrete signals only those with frequencies at or below one half of the sampling frequency will be properly represented after a transformation to the frequency domain. Therefore, to accurately determine the row spacing the image must be sampled at least twice as often as changes in rows occur. The spacing between tobacco rows can be roughly estimated as 80 pixels, so the maximum sample length to satisfy the Nyquist criterion would be 40 pixels. However, the sampling rate will also determine the resolution of the frequency domain, so a smaller sample length leads to increased frequency resolution, and therefore a more accurate estimation of the row spacing. All of the above considered, the initial sample length was selected as 10 pixels. The resulting data ensemble used to determine the row spacing consists of 10 data vectors, each 100 pixels wide in  $X$  and containing samples 10 pixels long in  $Y$ .

To create the data ensemble in software, a user is first asked to draw a bounding box around the desired area of a selected image. Figure 5.3 shows one image used in the study and the user-generated bounding box (blue). This allows the user to control what area of the image is sampled for frequency analysis. An important, but often overlooked factor when using the DFT, is that the transform assumes the data vector is periodic, and one entire period has been captured in the vector. By allowing a user to draw a bounding box on the image, a window is effectively being applied to the image data which then creates a roughly periodic data vector. After the bounding box has been selected, the data ensemble is generated sample by sample. Recall that each sample is a two-dimensional grouping of pixels (100 by 10), and the corresponding entry into the data ensemble is the mean pixel intensity of that sample. For the case of row detection, the mean pixel intensity is calculated in the usual way, as of interest are areas of bright pixels versus dark pixels. Figure 5.4, also shown below, presents a montage of two images; in the left image (a) only the first data vector has been generated, and on the right image (b) the entire data ensemble has been generated.



Figure 5.3: Tobacco image with user drawn bounding box (blue)

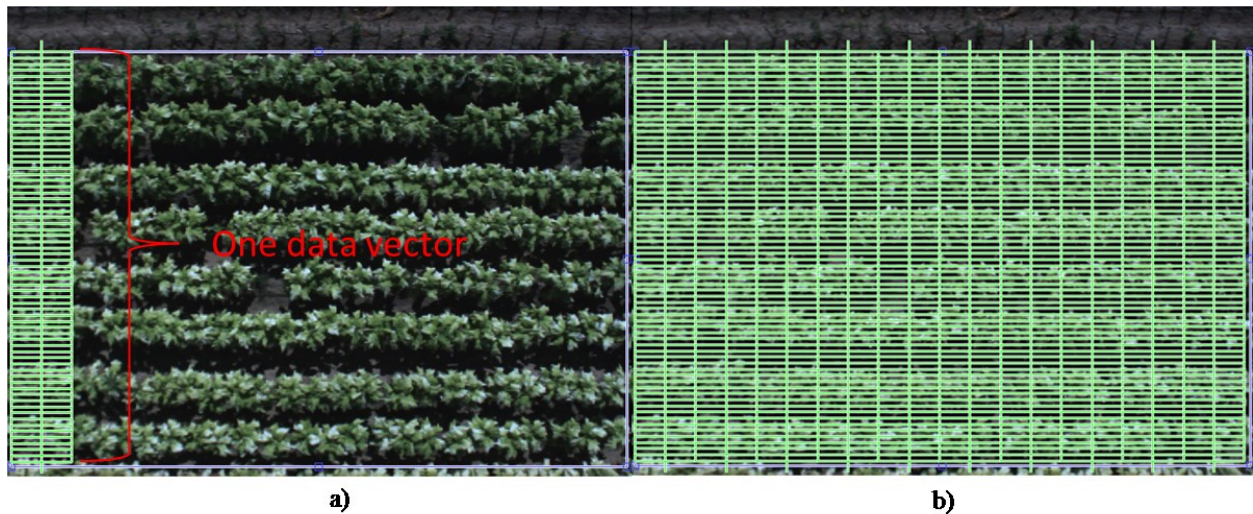


Figure 5.4: Generating data ensemble; a) a single data vector, b) the entire data ensemble.

With the data ensemble generated over the desired image area, the next step is to form the spectrum of each individual data vector. The spectrum of each data vector is computed in software with a Fast Fourier Transform (FFT) algorithm. The FFT makes use of the inherent symmetry of the DFT algorithm to speed up computation time. Figure 5.5 displays the resulting spectrum of the first data vector as pictured in Figure 5.4a. Even though a peak frequency is clearly apparent in the first spectrum, the remaining spectrums computed from the data ensemble are averaged together to improve the estimation of the row spacing frequency. The final averaged spectrum is shown in Figure 5.6. The frequency at which the maximum amplitude occurs (0.01194) is selected as the estimate for the row spacing frequency. By taking the reciprocal of the row spacing frequency, the resulting average row pitch is 84 pixels. Recall, the rows were initially estimated to be spaced about 80 pixels apart simply by inspecting Figure 5.2. The DFT method shows strong evidence that the average distance between rows is instead 84 pixels. There is likely variation in the row spacing across the entire field, and it is also likely that each row has not grown to the exact same size as the others. By creating a data ensemble that contains multiple vectors, each of which covers multiple rows, these slight variations are accounted for in the averaging of the spectrums, and the DFT results should be sufficient in bounding the area of each row.

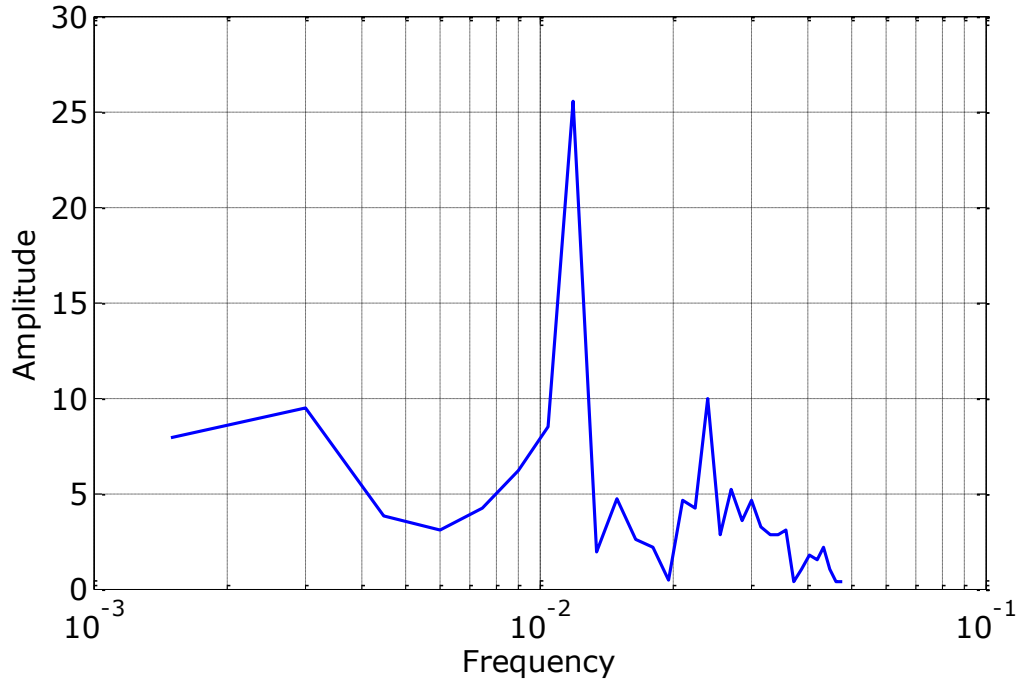


Figure 5.5: Single tobacco row spectrum




Figure 5.6: Averaged tobacco row spectrum

The final step before conducting the stress analysis is to create individual bounding boxes for each row. In this way, the desired row by row analysis can easily be completed. The simplest

method to creating the row boundaries is to divide the existing bounding box drawn on the image, like in Figure 5.3, into multiple bounding boxes each containing a single row. These new bounding boxes use the average row spacing estimate for division in the  $Y$  direction, and span the entire width of the image. Figure 5.7 shows 8 rows highlighted by their bounding boxes in another of the tobacco study images. With a boundary surrounding each row, it is possible to crop a row from the initial image, calculate the average NDVI of the individual row, and estimate which rows are under the greatest stress.




Figure 5.7: Tobacco image with row boundaries highlighted (red).

### 5.2.3 Stress Estimation

To begin the stress estimation process, the vegetation present in the tobacco study images needed to be segmented from the background noise. It was desired, once again, for the formation of an NDVI image to act as a single stage of segmentation, isolating the vegetation and blacking out soil and shadows. However, unlike with the flight images collected over the corn field, many of the shadows present in a sample tobacco study image have a slightly positive NDVI. Recall from the second chapter that this is not an unknown phenomenon when using NDVI. Depending on what is underneath the shadows, whether it is soil, water, or more vegetation, the shadowed area

could have a positive NDVI value if the NIR reflectance is greater than the Red reflectance. Unfortunately, this is the case for the tobacco study images. On average, shadowed areas have greater pixel intensity in the NIR image than in the visible image. The result being that when calculating the NDVI pixel by pixel, shadows end up with a slightly positive NDVI, and could consequently be misinterpreted as vegetation. The vast majority of shadows in the tobacco images lay over the soil between rows, so the best guess as to what is causing the undesirable behavior is the specific soil content at the Hudson Farms site. Bare soil, on the other hand, does not exhibit this behavior and consistently has a negative NDVI, resulting in proper elimination in the NDVI image.

To eliminate shadows from the NDVI image, an additional step is taken when performing the pixel by pixel calculations. By inspecting a visible tobacco study image in grayscale, it becomes apparent that the shadow covered soil has the lowest pixel intensity of any object in the scene, and can be reliably characterized by a pixel intensity below 30. It follows then, that the simplest solution is to not calculate the NDVI of pixels whose grayscale intensity in the visible image is below a selected threshold. Figure 5.8 shows two selections of an NDVI image formed in two ways; on the left (a) the NDVI image is formed in the usual way, and vegetation becomes obscured by the shadow (circled in red), on the right (b) the NDVI image is formed using the threshold technique, and the vegetation is now more clearly isolated. By adding a simple decision step to the previous NDVI algorithm, the tobacco can be segmented from the soil and the shadows in a single stage. The updated algorithm for forming the NDVI image is shown below in Algorithm 4, with shadow rejection added in line 8.




Figure 5.8: Tobacco sample segmented by two methods; a) usual NDVI method, b) NDVI method with shadow rejection. Notice how the shadow in a) (circled in red) has been eliminated in b)

**Algorithm 4** Form NDVI Image with Shadow Rejection

---

```

6.  $VIS \leftarrow$  read ‘visible image’
7.  $NIR \leftarrow$  read ‘near-infrared image’
8. Set Red equal to the first layer in VIS
9. Set Gray equal to GRAYSCALE (VIS)
10. Set grayscale Threshold for shadow rejection
11. for  $i = 1$  to image width
7.   for  $j = 1$  to image height
8.     if  $Gray(i, j) > Threshold$ 
9.       Form NDVI image by Equation 5.1            $\{(NIR - Red) ./ (NIR + Red)\}$ 
10.    else
11.      Set NDVI image( $i, j$ ) to 0            $\{\text{black in greyscale}\}$ 
12.    end if
13.  end for
14. end for
15. return  $\rightarrow$  NDVI image            $\{\text{greyscale image}\}$ 

```

---

Now that the tobacco has been successfully segmented, and each tobacco row individually bounded, it was a simple matter to calculate the MPI of each row and thusly the average NDVI. Figure 5.9 displays the same image as in Figure 5.7, transformed by the NDVI segmentation process and ready for row by row analysis. Inspection of Figure 5.9 reveals that very little of the vegetation remains after segmentation, and as a result the average NDVI of each row to be is expected to be low. Table 5.1 presents the average NDVI for each of the first 16 rows, and it is indeed the case that the averages are quite low. Recall the previous discussion in the second chapter of how late-season tobacco will see a significant decrease in NIR reflectance. Maturity of tobacco therefore leads to a drop in average NDVI. Also affecting the average NDVI of each row is the number of bleached leaves per stalk. The yellowing or bleaching of leaves is a visible sign of an unhealthy tobacco plant, and in practice it can be difficult to accurately diagnose the cause of the stress. Regardless of the cause, the bleached leaves exhibit a severe increase in Red reflectance over NIR reflectance and consequently come up negative on the NDVI scale. If the maturity of the crop is coupled with the large number of bleached leaves, it is not surprising that the rows of interest have a low average NDVI.




Figure 5.9: Sample tobacco study NDVI image

Table 5.1: Average NDVI for first 16 tobacco rows

Row	Average NDVI	Row	Average NDVI
1	0.252	9	0.201
2	0.258	10	0.201
3	0.240	11	0.185
4	0.199	12	0.194
5	0.187	13	0.175
6	0.191	14	0.187
7	0.172	15	0.200
8	0.171	16	0.214

The final step of the stress estimation process is to identify the block of most stressed rows. For this analysis a block of tobacco consists of 8 consecutive rows. The selection was made because additional chemical application to the tobacco was most likely to occur on an 8 row basis due to the machinery present at the farm. By using the average NDVI as an estimate for overall stress, a moving average was applied to the results shown in Table 5.1 to find the block with the lowest average. Table 5.2 displays the results of an eight row moving average computed over the first 16 rows. Two blocks (rows 6 – 13, and 7 – 14) have been identified to have the lowest average NDVI. The reported estimate then, is that blocks 6 and 7 are the most stressed rows present in



the area of interest. The precise cause of stress has yet to be evaluated, although likely causes are plant age, water stress due to flooding, and nutrient deficiency. Until more information regarding the crop status is made available, the NDVI analysis will have to suffice for overall stress estimation.

Table 5.2: Moving average stress estimation

<b>Block</b>	<b>Rows</b>	<b>Average NDVI</b>
1	1 – 8	0.209
2	2 – 9	0.203
3	3 – 10	0.195
4	4 – 11	0.188
5	5 – 12	0.188
6	6 – 13	0.186
7	7 – 14	0.186
8	8 – 15	0.189
9	9 – 16	0.195

#### 5.2.4 Bleached Leaf Segmentation

To conclude the late-season tobacco study a method for segmenting the bleached leaves from the green leaves was developed. One possible use for this segmentation would be to calculate a rough estimate of the percentage of bleached leaves to total tobacco vegetation. This estimation of bleached leaf concentration could then be used to identify areas of a tobacco field suffering pest or nutrient stress. For the bleached leaf segmentation to be useful in a wide variety of ways two results would need to be delivered: an image mask that isolates the healthy green tobacco from the bleached leaves, as well as background noise (soil, shadows, etc), and a second mask that isolates the bleached leaves from the healthy leaves and background noise. It is important to create two masks, so comparisons between bleached versus green can be made with minimal interference from background noise. A final consideration when developing the bleached leaf segmentation was to attempt to use only data from the visible images, so that in the future, a system would not necessarily need multispectral capabilities to detect bleached leaves in an otherwise healthy crop.

The first step to successful segmentation is to understand what features can be used to easily distinguish the content of interest from the background. Looking back to Figure 5.2, the image content can be divided into four primary groups: green leaves, bleached leaves, bare soil, and shadowed soil. Table 5.3 gives an average RGB pixel value for each of the four groups, as well as a primary feature in the visible spectrum that could be potentially used for segmentation. It should be noted that only one of the many possible RGB pixel values is listed for each of the groups, and for the leaf segmentation to be robust it must be able to account for this wide variation of pixel values. For example, the majority of bleached leaves are significantly brighter than any other content in the image, but there are still bleached leaves, distinguished by their yellow color, that have a similar brightness as green leaves. Therefore, pixel brightness alone cannot segment the green leaves from the bleached and does not account for the variety of bleached leaves in the scene. However, Table 5.3 still provides useful information that can be intelligently combined to achieve a successful segmentation of both the green and bleached tobacco.

Table 5.3: Typical tobacco image content

<b>Group</b>	<b>RGB Value</b>	<b>Distinguishing Feature</b>
Green Tobacco	94, 95, 67	Green pixel value significantly greater than blue value
Bleached Tobacco	148, 156, 166	Brightest average pixel values in image
Bare Soil	58, 55, 56	Low overall pixel intensity
Shadowed Soil	14, 16, 15	Low overall pixel intensity

Algorithm 5, as seen below, displays one possible implementation of the bleached leaf segmentation. The segmentation process begins with creating two blank masks the size of the visible tobacco image. Each mask is a binary image, with black pixels representing uninteresting content and white pixels representing either the green or bleached leaves respectively. The visible image supplied to the segmentation algorithm is searched pixel by pixel, and decisions are made to fill each mask with the corresponding content while rejecting background noise. The first condition used to segment the image is the total intensity of the given pixel in the green and blue layers. This first condition is used to distinguish the brightest pixels from the remainder of the image and identify them as bleached leaves. If the green and blue intensity is below a

selected threshold it moves on to the second condition. Here the primary objective is to determine whether the pixel represents a green plant or is part of the background. Recall from Table 5.3 that the most distinguishing feature of green leaves is that their green pixel values tend to be significantly greater than their blue pixel values. This is, of course, what gives the healthy leaves their green color in the visible image. Bleached leaves, bare soil, and shadowed soil all have pixels with roughly equal green and blue pixel values. It follows then, that the second segmentation condition looks at the difference between green pixel value and blue pixel value, where pixels with a difference above a second threshold are considered green leaves, and all others are considered background noise. Also, in the second segmentation condition, any green leaves that were sufficiently bright enough to be considered bleached in the first condition are removed from the bleached leaf mask. After each pixel has been evaluated by the algorithm, the two binary masks are returned and can then be used to isolate the leaves interest from the visible image, as seen in Figure 5.10.

---

**Algorithm 5** Bleached Leaf Segmentation
 

---

```

1.  $VIS \leftarrow$  read 'visible image'
2. Create blank binary image  $BleachedMask$  equal to  $SIZE(VIS)$            {image is entirely black}
3. Create blank binary image  $GreenMask$  equal to  $SIZE(VIS)$ 
4. Set  $ThresholdBright$  for identifying bright bleached leaves
5. Set  $ThresholdGreenLevel$  for identifying green leaves
6. for  $i = 1$  to image width
7.     for  $j = 1$  to image height
8.         if  $VIS(i, j, GREEN) + VIS(i, j, BLUE) > ThresholdBright$ 
9.             Set  $BleachedMask(i, j)$  to 1                               {white in image}
10.        elseif  $VIS(i, j, GREEN) - VIS(i, j, BLUE) > ThresholdGreenLevel$ 
11.            Set  $GreenMask(i, j)$  to 1
12.            Set  $BleachedMask(i, j)$  to 0                               {black in image}
13.        end if
14.    end for
15. end for
16. return  $\rightarrow BleachedMask, GreenMask$                          {binary images}

```

---




Figure 5.10: Bleached leaf segmentation example: a) original visible image, b) green leaves segmented from visible image, c) bleached leaves segmented from visible image

The masks created by the leaf segmentation process can be used for more than displaying an image of only the green leaves, or an image of only the bleached leaves. One possible use of the masks is to estimate the percentage of bleached leaves to total tobacco vegetation. If there is confidence that the pixels in each mask represent only vegetation of its respective group, then the masks can be combined and used to estimate the total amount of tobacco present in a scene. This is achieved by computing the number of non-zero pixels in the combined mask. Again, if there is confidence in the masks, then by following the same logic the amount of bleached tobacco in a scene can be estimated by calculating the number of non-zero pixels in the bleach leaf mask. A simple ratio can then be formed from each pixel count and the result used as an estimation of the percentage of bleached leaves per total tobacco captured in the image. If this method is applied to the images in Figure 5.10, the total tobacco vegetation count comes to 249990 pixels and the

bleached leaf count to 49362 pixels. The result is roughly 20% of the tobacco is bleached; meaning that one in every five leaves is likely to be bleached. Table 5.4 summarizes the bleached leaf percentage results described above. A simple inspection of the visual image gives confirmation that a significant portion of the tobacco contains bleached leaves, and provides a sanity check for the above estimation method. This method also helps to explain the low average NDVI per row results of Table 5.1. If one in every five leaves is bleached, it should have a negative NDVI value and significantly lower the average NDVI of its corresponding row. Overall, the bleached leaf segmentation has proved a success, and is able to provide a reasonable estimation of the bleached tobacco content per image.

Table 5.4: Bleached tobacco content estimation

<b>Tobacco Type</b>	<b>Pixel Count</b>
Green	200628
Bleached	49362
Total	249990
<b>Percent Bleached</b>	19.7%

# Chapter 6

## Summary & Conclusions

This thesis has presented the research, development, and testing of an unmanned aerial system for crop monitoring. The goal of the project was to investigate what potential capabilities the system has in the wide field of agronomic remote sensing. Specifically, what useful results could be generated from multispectral images and how could those results be interpreted to inform farmers about the status of their crop. The system developed for the project integrates a multispectral camera payload to an autonomous helicopter to form the complete UAS. Covered in the third chapter are details about the integration of the payload to the vehicle, and the basic operation of the system during an unmanned flight operation. Images in both the visible and near-infrared spectrums are gathered simultaneously and post-processed to isolate and analyze the vegetation. Crops studied over the duration of this project were corn grown at Kentland Farms near Blacksburg, VA and late-season tobacco grown in South Hill, VA. Autonomous flight missions took place over both crops, yielding large image data sets. Each crop was growing under different conditions and provided various ways in which to test the abilities of the system, and inspire the development of different analysis techniques. This work has focused on the processing of those images collected from the multispectral payload, both on the ground and in the air, and the first section of this chapter will summarize the results of the image analysis as they apply to potential crop monitoring capabilities. The second section of this chapter will present recommendations for future work that could be used to provide additional capabilities to the UAS and expand the crop monitoring project as a whole. The final section presents the

authors concluding remarks regarding the work completed and the impact of UAS on the agricultural practices of the future.

## 6.1 Summary of System Capabilities

As previously stated, the primary goal of this work was to explore and develop capabilities of a UAS for crop monitoring. The literature review presented in the second chapter discussed many possible applications for agronomic remote sensing, both as applied today and as desired for the future. A topic common in many of the works presented in the review was the use of NDVI for segmenting vegetation from background noise in a scene. However, many found NDVI to not be entirely effective as a segmentation tool due to its sensitivity to soil interference. This work was able to overcome that obstacle by creating NDVI images and using simple machine vision techniques to remove soil and other noise from the images. Specifically, the techniques used were thresholding, to remove shadows from the NDVI image, and binary mask application, to reduce soil effects. The ability for the multispectral system to properly isolate the vegetation of interest using NDVI images leads to increased estimation fidelity when those images are analyzed further.

Another topic that appeared many times in the literature review was estimation of nutrient status, particularly nitrogen nutrient status, in order in to develop better practices for applying the fertilizer and reducing nitrate runoff. This topic was addressed specifically with data collected during a flight test over the corn field. Least Significant Difference statistical tests were used to assess if the multispectral system could reliably distinguish corn of different fertilizer levels. This was possible thanks to the controlled nature of the fertilizer rate study taking place in the field and the ground truth data provided by Dr. Wade Thomason. The LSD tests showed that corn whose nitrogen fertilizer rates differed by 100 lb/ac or more can reliably be detected by the mean pixel intensity of a properly formed NDVI image. Test results for the HR hybrid also showed the ability to detect when a corn crop has reached its optimum nitrogen uptake. What this means is that the UAS could be used in field-wide studies to detect crops whose nutrient level is significantly different than its neighbors, and those crops that are no longer making use of additional fertilizer. The capability then that has been developed is one for nutrient status estimation and monitoring. As mentioned in the literature review, if a reasonable estimate of a

crop's nitrogen status can be made, then fertilizer can be applied "as-needed" in the form of a sidedress. The "as-needed" nitrogen management schemes have been proven to be much more effective in applying appropriate amounts of fertilizer to encourage maximum yield while limiting nitrate runoff. The capability for the unmanned system to monitor a crop's nitrogen status is one of the most impactful of this work, and represents a success for UAS in potentially improving agricultural practices.

Although not specifically addressed as a key issue in crop research, the ability to detect the row spacing in a field demonstrated in this work has many practical uses. In particular, by estimating the row spacing, row by row analysis of a crop becomes possible with minimal human interaction. An accurate estimation of row spacing in a field becomes important when the region of interest is at the row level as opposed to the general field-wide level, or very precise plant or leaf level. As agricultural practices move forward, from macro based decision making to true micro-level management, it is likely that the transition will begin with decisions made at the row level. Therefore, the simple row spacing estimation method provided in this work will prove useful for continuing the use of imaging systems in agronomic remote sensing.

The last capability developed thus far in the project is the bleached tobacco leaf segmentation presented in Chapter 5. The number of bleached leaves in a tobacco crop says a lot about that crop's status. By isolating the bleached and green leaves, both from each other and background noise, it is possible to analyze them separately for parameter estimation. On the other hand, if there is strong confidence in the segmentation results, it is possible to estimate the concentration of bleached leaves to total tobacco in a given image. When performed on a test image from the late-season tobacco study, an estimated 20% of the tobacco leaves were bleached. This was a considerable amount and helped to explain the high level of stress present in the region of interest. The same bleached leaf concentration estimation could have been performed field-wide and acted as another factor for monitoring crop status. The bleached leaf segmentation was possible with only the visible image, which would allow a simpler system to perform the same analysis without relying on a more expensive multispectral sensor. It is the opinion of the author that estimating the bleached leaf concentration is the most impactful result of the late-season tobacco study, and could be applied in the future to detect areas of nutrient or pest stress, as well as improve the timing for harvest.



## 6.2 Future Works

This work finally comes to a close with a look to the future of the crop monitoring project at the Unmanned Systems Lab. It is the hope of the author, that it has been made clear just how great the potential is for unmanned systems to revolutionize agriculture. That being said, there are countless avenues to pursue as the project moves forward. Presented here are just a few out of the many directions to be explored, but with a focus on those that have been determined as most beneficial to sustaining the crop monitoring project for the long term.

Firstly, it is desired to develop the capability to form field-wide image mosaics. These mosaics are single images, very large in size, that have been created by stitching together the individual images gathered during a test flight. Mosaics can provide a method of feedback to a farmer that is as simple as a color-coded map of his field. Although the current system can report the field-wide status of a crop, it does so only in graphs and tables. If however, this was possible with a map, the status of the crop might be more easily understood and appropriate action taken with minimal delay. It cannot be stressed enough the advantage of using a visual medium to communicate results over that of numbers and words. Software already exists that can take large image sets, complete the stitching, and return a geo-referenced mosaic. Unfortunately, most commercial mosaic software is in a price range beyond the limit of the project. It is the author's hope then, to partner with a software developer or for custom software to be written. In any case, a geo-referenced mosaic provides an extremely helpful method for communicating the status of a field and will be essential for the micro-level management practices of the future.

Secondly, it is recommended to acquire additional sensors that allow for narrow-band multispectral imagery. This sensor could be in the form of a hyperspectral camera, or perhaps, a suite of cameras with narrow-band filters. As previously mentioned, significant research efforts have been undergone to understand specific reflectance characteristics of many crops, and specialized vegetation indexes created to identify particular crop parameters. Currently, the multispectral system of this work relies on a broad-band NDVI calculation, and although significant results were achieved, research has shown that narrow-band indexes will improve the accuracy of crop stress and status estimations. In the same way that the system should expand its sensing capabilities, the amount of factors used in crop status estimation should be expanded to

include popular agronomic inputs such as: average rain fall, soil content, and growing degree days. Data fusion techniques should be used to incorporate these additional factors, which could be provided by farmers on the ground, with the spectral data gathered by the new imaging sensors to improve parameter and stress estimations. Crops are a complex system, so by increasing the number of inputs to an estimation model, the fidelity of the results will in turn see an increase. Finally, along the lines of expansion, the crop monitoring UAS should be tested over a much wider variety of crop types to make the most out of the improved sensing abilities. Currently, the system has only been tested with two crops, and to extend the longevity of the system it must be able to monitor many different crops which might all be growing on the same farm, or in the same region.

Recall in Chapter 4, it was found that the pixel intensity of NDVI images is sensitive to changes in camera exposure over 250  $\mu$ s. In the future, it is recommended to explore in more depth what factors the multispectral system is sensitive to in order to obtain results for a sufficient analysis. These factors could include ambient lighting conditions, flight mission altitude, and soil content, to name a few. For the UAS to be viable in long term, it must retain mission flexibility and provide repeatable results. After understanding what factors the system is most sensitive to, a pre-flight camera calibration routine should be established. It would be ideal for the payload to make measurements itself and determine its own settings such as exposure time, focus, and trigger rate. In any event, whether the data is gathered by the system itself or supplied to the system by a user, a calibration routine will lend robustness to the system and help to ensure that data collected from any flight is useful and suitable for crop status analysis.

As is stands, field-wide data was only gathered at one particular time in the growing season for each of the crops studied in this work. Crops see significant changes in physiology and spectral reflectance over the duration of the season, and therefore it is recommended to apply the multispectral UAS to full season crop monitoring. Data collection should occur from planting through harvest, with the goal of developing models which could be inspected at a particular growth stage to compare the relative plant health. These models could be improved year after year as the data set grows and more factors are used to improve estimation fidelity. Again, in order to sustain the life of the project, it is critical that full season monitoring become the trend as opposed to intermittent data collection. It would be a waste of the system to continue to only

fly once a season, as one of the prime advantages of UAS is their ability to collect data with high temporal frequency.

Finally, the last suggestion for future work is to develop an autonomous ground robot to act in conjunction with the UAS. A ground based system could serve to collect the additional data, such as a soil analysis, and also be used for precise chemical applications. In the cooperative mission, the UAS acts as a field-wide source for data collection and relays that information to one or more ground robots capable of the more precise data collection and crop management. The ground robot then becomes the tool for “as needed” chemical applications, with the potential to significantly reduce soil damage and runoff. The use of multiple autonomous systems working in tandem would allow for sophisticated “smart field” systems and field-wide micro-level management practices to become a reality.

### **6.3 Concluding Remarks**

Overall, the unmanned aerial system presented in this work has both current capabilities and vast potential for monitoring field-wide crop status. It has been demonstrated that the multispectral images gathered by the system can be used for a variety of tasks such as vegetation characterization, nutrient deficiency detection, and grain yield estimations. Given access to new sensors and more data, the system has the potential to identify crop specific diseases, develop full season growth models, and provide the field-wide measurements necessary for "as-needed" agrochemical application schemes. Unmanned aerial vehicles make the ideal platforms to gather all of the data required for the agricultural practices of the future. Their low altitude flights generate data at the high resolutions needed for accurate crop parameter estimations, and their ability to fly missions at high temporal frequencies ensure that timely feedback for micro-level management techniques is present. It is the opinion of the author, that UAS will continue to be instrumental in advancing the agriculture industry, and that the system presented within is just one step towards more intelligent, productive, and responsible farming.

# Bibliography

- [1] M. A. Sutton, O. Oenema, J. W. Erisman, A. Leip, H. van Grinsen, and W. Winiwarter, "Too much of a good thing," *Nature*, vol. 472, pp. 159-161, 2011.
- [2] J. Weekley, J. Gabbard, and J. Nowak, "Micro-Level Management of Agricultural Inputs: Emerging Approaches." *Agronomy*, vol. 2, no. 4, pp. 321-357, 2012.
- [3] S. L. Osborne, J. S. Schepers, D. D. Francis, and M. R. Schlemmer, "Use of spectral radiance to estimate in-season biomass and grain yield in nitrogen-and water-stressed corn," *Crop Science*, vol. 42, no. 1, pp. 165-171, 2002.
- [4] M. W. Barret, E. K. Leggett, "The development of aerial infrared photography to detect *Echinochloa* species in rice," *Proceedings of the Seventh Asian-Pacific Weed Science Society Conference*, pp. 41-44, 1979.
- [5] C. J. Tucker, B. N. Holben, J. H. Elgin, and J. McMurtrey, "Relationship of spectral data to grain-yield variation," *Photogrammetric Engineering & Remote Sensing*, vol. 46, pp. 657-666, 1980.
- [6] T. J. Malthus and A. C. Madeira, "High resolution spectroradiometry: Spectral reflectance of field bean leaves infected by *Botrytis Fabae*," *Remote Sens. Environ.*, vol. 45, no. 1, pp. 107-116, Jul. 1993.
- [7] LARS, Laboratory for Applications of Remote-Sensing, "Remote Multi-spectral Sensing in Agriculture," Annual Report No. 3, Purdue University, West Lafayette, USA. 1968.
- [8] LARS, Laboratory for Applications of Remote-Sensing, "Multi-spectral Sensing in Agriculture," Annual Report No. 4, Purdue University, West Lafayette, USA, 1970.
- [9] Darrel L. Williams, Samuel Goward, and Terry Arvidson, "Landsat: Yesterday, Today, and Tomorrow," *Photogrammetric Engineering & Remote Sensing*, vol. 72, no. 10, pp. 1171-1178, 2006.
- [10] Kamlesh Lulla, M. Duane. Nellis, and Brad Rundquist, "The Landsat 8 is ready for geospatial science and technology researchers and practitioners," *Geocarto International*, vol. 28, no. 3, pp. 191, 2013.
- [11] Zhiyan Zhou, et al, "Canopy hyperspectral reflectance feature of rice caused by Brown Plant-hopper (*Nilaparvata lugens*) infestation," 2010 *ASABE Annual International Meeting*, June 2010.

- [12] Zhang, Minghua, Zhihao Qin, and Xue Liu, "Remote sensed spectral imagery to detect late blight in field tomatoes," *Precision Agriculture*, vol. 6, pp. 489-508, 2005.
- [13] D. Moshou, et al, "Simultaneous identification of plant stresses and diseases in arable crops using proximal optical sensing and self-organizing maps," *Precision Agriculture*, vol. 7, pp. 149-164, 2006.
- [14] J. W. Rouse, R. H. Haas, J. A. Schell, D. W. Deering, and J. C. Harlan, "Monitoring the vernal advancements and retrogradation of natural vegetation," NASA/GSFC, Final Report, Greenbelt, MD, p. 371, 1974.
- [15] David M. Gates, *Biophysical Ecology*, Springer-Verlag, New York, p. 611, 1980.
- [16] C. J. Tucker, "Asymptotic nature of grass canopy spectral reflectance," *Applied Optics*, vol. 16, no. 5, pp. 1151-1156, 1977.
- [17] A. R. Huete, "A soil-adjusted vegetation index (SAVI)," *Remote Sensing of Environment*, vol. 25, pp. 295-309, 1988.
- [18] J. Penuelas, I. Filella, C. Biel, L. Serrano, and R. Save, "The reflectance at the 950-970 nm region as an indicator of plant water status," *International Journal of Remote Sensing*, vol. 14, pp. 1887-1905, 1993.
- [19] Ranga B. Myneni, R. R. Nemani, and Steven W. Running, "Estimation of Global Leaf Area Index and Absorbed Par using Radiative Transfer Methods," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 35, no. 6, 1997.
- [20] J. Gamon, et al, "Relationships between NDVI, canopy structure, and photosynthesis in three Californian vegetation types," *Ecological Applications*, vol. 5, pp. 28-41, 1995.
- [21] K. C. Swain, H. P. W. Jayasuriya, and V. M. Salokhe, "Low Altitude Remote Sensing (LARS): A Potential Substitution to Satellite Based Remotes Sensing for Precision Agriculture Adoption in Fragmented and Diversified Farming Conditions," *Agricultural Engineering International: the CIGR Ejournal*, Invited Overview No. 12, Vol. IX, Sept. 2007.
- [22] Sugiura, Ryo, Noboru Noguchi, and Kazunobu Ishii. "Remote-sensing technology for vegetation monitoring using an unmanned helicopter," *Biosystems Engineering* 90.4, pp. 369-379, 2005.
- [23] Jose A. J. Berni, et al, "Thermal and Narrowband Multispectral Remote Sensing for Vegetation Monitoring From an Unmanned Aerial Vehicle," *IEEE Transactions on Geoscience and Remote Sensing*, 2009.

- [24] W. C. Bausch, and C. M. U. Neale, "Crop coefficients derived from reflected canopy radiation: a concept," *Transcriptions from ASAE* 30.3, pp. 703-709, 1987.
- [25] W. C. Bausch, "Soil Background Effects on Reflectance-Based Crop Coefficients for Corn," *Remote Sens. Environ.*, vol. 46, pp. 213-222, 1993.
- [26] Chang, Jiyul, et al, "Corn (L.) Yield Prediction Using Multispectral and Multidate Reflectance," *Agronomy Journal* 95.6, pp. 1447-1453, 2003.
- [27] R. H. Fox, J. M. Kern, and W. P. Piekielek, "Nitrogen fertilizer source, and method and time of application effects on no-till corn yields and nitrogen uptake," *Agronomy Journal* 78, pp. 741-746, 1986.
- [28] W. C. Bausch, and H. R. Duke, "Remote Sensing of Plant Nitrogen Status in Corn," *Transcriptions from ASAE* 39.5, pp. 1869-1875, 1995.
- [29] S. L. Osborne, J. S. Schepers, D. D. Francis, and M. R. Schlemmer, "Detection of Phosphorus and Nitrogen Deficiencies in Corn Using Spectral Radiance Measurements," *Agronomy Journal* 94.6, pp. 1215-1221, 2002.
- [30] H. Lichtenthaler, A. Gitelson, and M. Lang, "Non-destructive determination of chlorophyll concentration of leaves of a green and an aurea mutant of tobacco by reflectance measurements," *J. Plant Physiol.*, vol. 148, pp. 483-493, 1996.
- [31] Anatoly A. Gitelson, and Mark N. Merzlyak, "Remote sensing of chlorophyll concentration in higher plant leaves," *Adv. Space Res.*, vol. 22, no. 5, pp. 689-692, 1998.
- [32] Hongbo Qiao, et al, "Study on Relationship between Tobacco Canopy Spectra and LAI," *International Federation for Information Processing AICT* 345, pp. 650-657, 2011.
- [33] Ezekia Svotwa, et al, "Selection of Optimum Vegetative Indices for the Assessment of Tobacco Float Seedlings Response to Fertilizer Management," *ISRN Agronomy*, vol. 20, 2012.
- [34] F. Esposito, R. Giancarlo, and Antonio Moccia. "1st mini-UAV integrated hyperspectral/thermal electro-optical payload for forest fire risk management," *Proc. AIAA Infotech Aerosp. Conf.*, vol. 1, pp. 653-665, 2007.
- [35] JAI Ltd., "User's Manual AD-080GE," Ver. 1.1, Shanghai, China, June, 2012.
- [36] D. Freiman, Z. Elgart, and Arie Chernoborov. "Dichroic beamsplitter for a dual-wavelength camera in a long-range airborne reconnaissance system," *Proc. SPIE.*, vol. 5783, pp. 827-834, 2005.

- [37] J. Gamon, J. Penuelas, and C. Field, "A narrow-waveband spectral index that tracks diurnal changes in photosynthetic efficiency," *Remote Sens. Environ.*, vol. 41, no. 1, pp. 35–44, 1992.
- [38] "fit-PC2 Features," 2013. URL: <http://www.fit-pc.com/web/products/fit-pc2/>
- [39] Prasad S. Thenkabail, Ronald B. Smith, and Eddy De Pauw, "Evaluation of Narrowband and Broadband Vegetation Indices for Determining Optimal Hyperspectral Wavebands for Agricultural Crop Characterization," *Photogrammetric Engineering & Remote Sensing*, vol. 68, no. 6, pp. 607–621, 2002.
- [40] B. E. Bayer, "Color Imaging Array," United States Patent 3,971,065, July 20, 1976.
- [41] Sindhuja Sankaran, Ashish Mishra, Reza Ehsani, and Cristina Davis, "A review of advanced techniques for detecting plant diseases," *Computers and Electronics in Agriculture*, vol. 72, pp. 1-13, 2010.
- [42] D. W. Lamb, "The use of qualitative airborne multispectral imaging for managing agricultural crops — a case study in south-eastern Australia," *Australian Journal of Experimental Agriculture*, vol. 40, pp. 725-738, 2000.
- [43] R. G. D. Steel, and J. H. Torrie, *Principles and procedures of statistics*, McGraw-Hill Inc., New York, 1960.
- [44] Lyman Ott, Michael Longnecker, *An introduction to statistical methods and data analysis*, Brooks/Cole Cengage Learning, Belmont CA, 2010.
- [45] "Nitrogen Use Efficiency," DuPont Pioneer, 2013. URL: [http://www.pioneer.com/CMRoot/Pioneer/About\\_Global/our\\_research/pipeline/spec\\_sheets/spec\\_sheets\\_files/Corn\\_NUE\\_2013.pdf](http://www.pioneer.com/CMRoot/Pioneer/About_Global/our_research/pipeline/spec_sheets/spec_sheets_files/Corn_NUE_2013.pdf)
- [46] R. N. Bracewell, *The Fourier transform and its applications*, McGraw-Hill Inc., New York, 1986.
- [47] Arthur R. Weeks, *Fundamentals of electronic image processing*, Bellingham: SPIE Optical Engineering Press, 1996.