

A Prototype Polarimetric Camera for Unmanned Ground Vehicles

Mark Umansky

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

Alfred L. Wicks, Committee Chair

Kathleen Meehan

Dennis W. Hong

June 28, 2013

Blacksburg, VA

Keywords: Unmanned Ground Vehicles, Robotic Perception, Polarization, Machine Vision

A Prototype Polarimetric Camera for Unmanned Ground Vehicles

Mark Umansky

Abstract

Unmanned ground vehicles are increasingly employing a combination of active sensors such as LIDAR with passive sensors like cameras to perform at all levels of perception, which includes detection, recognition and classification. Typical cameras measure the intensity of light at a variety of different wavelengths to classify objects in different areas of an image. A polarimetric camera not only measures intensity of light, but can also determine its state of polarization.

The polarization of light is the angle the electric field of the wave of light takes as it travels. A polarimetric camera can identify the state of polarization of the light, which can be used to segment highly polarizing areas in a natural environment, such the surface of water. The polarimetric camera designed and built for this thesis was created with low cost in mind, as commercial polarimetric cameras are very expensive. It uses multiple beam splitters to split incoming light into four machine vision cameras. In front of each machine vision camera is a linear polarizing filter that is set to a specific orientation. Using the data from each camera, the Stokes vector can be calculated on a pixel by pixel basis to determine what areas of the image are more polarized. Test images of various scenes that included running water, standing water, mud, and vehicles showed promise in using polarization data to highlight and identify areas of interest. This data could be used by a UGV to make more informed decisions in an autonomous navigation mode.

Acknowledgements

I am very lucky to have had the opportunity to study at one of the world's best universities for robotics research for the past six years. I have learned more than I could have ever imagined when I started.

I would like to thank my advisors, Dr. Wicks, Dr. Meehan, and Dr. Hong for their help and guidance in the final years of my undergraduate degree and throughout my graduate studies.

I would also like to thank my lab mates and colleagues, especially Ben Goldman, William French, and Michael Teresi for their help and experience in going through graduate school together. I also want to thank Dr. John Bird and Dr. Jonathan Gaines for their guidance and the initial opportunity to pursue robotics research outside of the classroom that made my graduate education possible.

I would like to thank my parents and my brother for the lifetime of love and support. Finally I want to thank Laura for sticking with me throughout the years and for the much needed editing my writing required.

Table of Contents

Abstract.....	iii
Acknowledgements.....	iii
Table of Contents.....	iv
List of Acronyms and Definitions.....	vii
Table of Figures.....	viii
List of Tables.....	xi
1 Background and Motivation.....	1
1.0 Thesis Summary.....	1
1.1 UGVs at Virginia Tech.....	1
1.1.1 2004 DARPA Grand Challenge.....	2
1.1.2 2005 DARPA Grand Challenge.....	2
1.1.3 2007 DARPA Urban Challenge.....	3
1.1.4 GUSS Program.....	5
1.2 UGV Sensing Methods.....	6
1.2.1 Active Sensing.....	6
1.2.2 Passive Sensing.....	7
1.3 Why Polarized Light?.....	9
1.4 Thesis Content.....	10
2 Literature Review.....	12
2.0 Chapter Summary.....	12
2.1 Discovery of Polarization.....	12
2.2 Polarization in nature.....	12
2.2.1 Honeybee orientation with skylight polarization.....	12
2.2.2 Aquatic insects and water detection.....	13
2.2.3 Human Polarization Sensitivity.....	13
2.3 Polarization Cameras.....	14
2.3.1 Division of Time polarization camera.....	14
2.3.2 Polarization camera using a beam splitter.....	14
2.3.3 Polarization camera using a beam splitter and waveplates.....	16
2.4 Scene Classification and Image Processing.....	17
2.4.1 Mud detection using polarization.....	17

2.4.2	Water detection using passive sensing.....	18
2.4.3	Haze Reduction.....	19
2.5	Thesis Relevance	21
3	Methods of Camera Design and Assembly.....	22
3.0	Introduction.....	22
3.1	Fundamental Science and Optics Hardware	22
3.1.1	Electromagnetic Waves and Polarized Light	22
3.1.2	Jones Vector.....	24
3.1.3	Stokes Parameters	25
3.1.4	Jones and Mueller Matrices	27
3.1.5	Brewster’s Angle.....	28
3.1.6	Polarizing Filters	30
3.1.7	Beam splitters.....	31
3.2	Detailed design and assembly	32
3.2.1	Requirements	32
3.2.2	Design Summary.....	33
3.2.3	Image Registration	33
3.2.4	Optical Components and Camera Details	35
3.2.5	Base Platform.....	39
3.2.6	Beam Splitter Layout	39
3.2.7	Beam Splitter Mount Design and Results	41
3.2.8	Polarizing Filters and Holder	42
3.2.9	Software	46
4	Experiments	49
4.0	Summary of Chapter	49
4.1	Indoor verification and tests.....	49
4.1.1	Image alignment between Firefly cameras	49
4.1.2	Intensity differences between cameras.....	53
4.1.3	Stokes images of a polarized light source	57
4.2	Polarization in the outdoors	60
4.2.1	Mud and Still Water.....	60
4.2.2	Moving water	63
4.2.3	Man-made objects.....	67

5	Results.....	71
5.1	Camera Alignment.....	71
5.2	Artificially Polarized Light Source.....	71
5.3	Water and Mud.....	71
5.4	Man Made Objects.....	72
6	Conclusion.....	73
6.0	Summary of Chapter.....	73
6.1	Polarimetric Camera Design and Assembly.....	73
6.1.1	Overall Cost.....	74
6.2	Future work.....	74
6.2.1	Redesign of beam splitter platform.....	74
6.2.2	Redesign of filter holder.....	75
6.2.3	Alternative camera and lens selection.....	76
6.2.4	Polarization data sets and other polarizing sources.....	76
6.2.5	Simplified design.....	76
	References.....	77
	APPENDIX A: Typical Stokes and Jones vectors for polarized light.....	81
	APPENDIX B: Typical Jones and Mueller Matrices.....	82
	APPENDIX C: MSP430 Source Code.....	83
	APPENDIX D: Camera capture source code.....	84
	APPENDIX E: Procedure for using the polarimetric camera.....	85

List of Acronyms and Definitions

Acronym	Definition
FOB	Forward Operating Base
GUI	Graphical User Interface
LIDAR	Light Detection and Ranging
LWIR	Long Wave Infrared
NDVI	Normalized Difference Vegetation Index
NIR	Near-infrared
NSWC	Naval Surface Warfare Center
PCB	Printed Circuit Board
SWIR	Short Wave Infrared
UGV	Unmanned Ground Vehicle

Table of Figures

FIGURE 1: VIRGINIA TECH’S TWO ENTRIES INTO THE 2005 DARPA GRAND CHALLENGE, CLIFF (LEFT) AND ROCKY (RIGHT) [4].	2
FIGURE 2: THE VICTOR TANGO FORD ESCAPE HYBRID, ODIN, WITH SENSORS LABELED [5].	4
FIGURE 3: THE SENSOR COVERAGE AROUND ODIN [5].	4
FIGURE 4: THREE GUSS UNMANNED GROUND VEHICLES [6].	5
FIGURE 5: A VELODYNE 64-E LIDAR SCAN WITH OBJECTS SEGMENTED AS DIFFERENT COLORS [8].	6
FIGURE 6: A SET OF THREE SCENES TAKEN WITH VISIBLE (LEFT) AND NEAR-INFRARED (CENTER) SENSORS. THE NDVI SEGMENTED IMAGE IS ON THE RIGHT.	8
FIGURE 7: AN OUTDOOR SCENE TAKEN WITH A NEAR-INFRARED CAMERA (LEFT) AND THE CORRESPONDING DEPTH MAP (RIGHT).	9
FIGURE 8: A VISIBLE IMAGE OF THE DUCK POND AT VIRGINIA TECH.	10
FIGURE 9: A THRESHOLDED DEGREE OF POLARIZATION OF THE SAME SCENE FROM FIGURE 8.	10
FIGURE 10: DEGREE OF LINEAR POLARIZATION P OF LIGHT REFLECTED AT THE BREWSTER ANGLE FROM A POND WITH CLEAR WATER AND DARK BOTTOM (DASHED LINE) AND A EUTROPHIC POND WITH GREEN-BROWN APPEARANCE (SOLID LINE) [17].	13
FIGURE 11: POLARIZATION MODULE PRINCIPLE OF THE SALSA CAMERA [18].	14
FIGURE 12: TOP-VIEW DIAGRAM OF A TWO-CCD POLARIZATION CAMERA USING A POLARIZING BEAM SPLITTER [19].	15
FIGURE 13: SCHEMATIC VIEW OF THE FLUXDATA 3-CCD CAMERA SHOWING THE 3-WAY BEAM SPLITTER AND POLARIZER CONFIGURATION [21].	16
FIGURE 14: SCHEMATIC OF A BEAM SPLITTER POLARIZATION CAMERA THAT USES WAVE-PLATES [23].	16
FIGURE 15: A SET OF POLARIZATION CONTRAST AND SEGMENTED IMAGES OF A PATCH OF MUD TAKEN SEVEN HOURS APART [24].	18
FIGURE 16: (A) A RAW IMAGE TAKEN AT ONE POLARIZATION STATE. (B) A SECOND RAW IMAGE TAKEN AT A DIFFERENT POLARIZATION STATE. (C) THE FINAL PROCESSED IMAGE WITH HAZE REDUCED [29].	20
FIGURE 17: AN ILLUSTRATION OF AN ELECTROMAGNETIC WAVE [30].	22
FIGURE 18: OSCILLATION IN A WAVE OF LIGHT VIEWED END-ON. (A) UNPOLARIZED LIGHT, ALL PLANES ARE EQUALLY PROBABLE. (B) ONE WAVE POLARIZED LIGHT [31].	23
FIGURE 19: (LEFT TO RIGHT) REPRESENTATIONS OF LINEAR POLARIZATION, CIRCULAR POLARIZATION, AND ELLIPTICAL POLARIZATION. THE BOTTOM PLANE SHOWS THE SHAPE TRACED OUT BY THE COMBINATION OF THE TWO COMPONENTS, THE RED AND BLUE WAVES.	24
FIGURE 20: A VISUALIZATION OF P-POLARIZED AND S-POLARIZED LIGHT [32].	24
FIGURE 21: AN ILLUSTRATION OF THE POLARIZATION OF LIGHT WHEN REFLECTED OFF A DIELECTRIC SURFACE AT BREWSTER’S ANGLE [34].	29
FIGURE 22: THE REFLECTIVITY OF SMOOTH WATER AT 20° C [36].	30
FIGURE 23: (TOP) A ELECTROMAGNETIC WAVE POLARIZED PARALLEL TO A DICHROIC CRYSTAL IS ABSORBED AND ITS INTENSITY REDUCED. (BOTTOM) AN ELECTROMAGNETIC WAVE POLARIZED PERPENDICULARLY TO THE CRYSTAL PASSES THROUGH WITH LITTLE CHANGE. [38].	31
FIGURE 24: AN ARTIST REPRESENTATION OF A CUBE BEAM SPLITTER. THE IMAGE IS TRANSMITTED FROM THE LEFT SIDE, THROUGH THE BEAM SPLITTER AND REFLECTED AT A 90° ANGLE, PRODUCING TWO IMAGES [39].	32
FIGURE 25: A CAD SCREENSHOT OF THE OVERALL SYSTEM DESIGN SHOWING THE CAMERAS, BEAM SPLITTER ASSEMBLY, MOUNTING PLATE, AND USB HUB.	33
FIGURE 26: (TOP) A PAIR OF IMAGES SHOWING A ROOFTOP SCENE VIEWED AT DIFFERENT ANGLES. (BOTTOM) THE SAME TWO IMAGES WITH MATCHING SIFT FEATURES [41].	34
FIGURE 27: THORLABS BS016 50:50 NON-POLARIZING BEAM SPLITTING CUBE. [42]	36

FIGURE 28: THE POINT GREY FIREFLY MV MACHINE VISION CAMERA [43].	37
FIGURE 29: POINT GREY FIREFLY MV EXTERNAL TRIGGER TIMING CHARACTERISTICS [43].	38
FIGURE 30: A TOP DOWN VIEW OF THE MSP430 LAUNCHPAD DEVELOPMENT BOARD [44].	38
FIGURE 31: A CAD SCREENSHOT OF THE ALUMINUM BASE PLATE.	39
FIGURE 32: A TOP-DOWN VIEW OF THE BEAM SPLITTERS, LABELED 1 THROUGH 3.	40
FIGURE 33: A TOP-DOWN VIEW OF THE PATH LIGHT TAKES INSIDE THE BEAM SPLITTER ASSEMBLY. THE DARK BLUE ARROW IS THE ENTRANCE FACE OF THE ASSEMBLY. THE RED ARROWS SHOW HOW THE LIGHT IS SPLIT INSIDE. THE LIGHT BLUE LINES AND SQUARE ARE WHERE THE FIREFLY MVS ARE LOCATED.	40
FIGURE 34: A FRONT VIEW OF THE BEAM SPLITTER ASSEMBLY (NOT VISIBLE: BEAM SPLITTER #2).	40
FIGURE 35: A CAD SCREENSHOT OF THE BEAM SPLITTER CLAMP METHOD (HIGHLIGHTED IN BLUE).	41
FIGURE 36: AN IMAGE OF THE BEAM SPLITTERS DURING ASSEMBLY OF THE CAMERA.	42
FIGURE 37: A CAD SCREENSHOT OF THE FILTER HOLDER BOX. THE FIRST OPENING IS ON THE BOTTOM RIGHT OF THE BOX.	42
FIGURE 38: AN IMAGE OF THE FILTER HOLDER AFTER IT WAS PRINTED.	43
FIGURE 39: A DIAGRAM SHOWING HOW LIGHT FROM AN LCD SCREEN WITH -45° POLARIZATION IS FILTERED BY LINEAR POLARIZERS PLACED AT DIFFERENT ANGLES.	44
FIGURE 40: A HISTOGRAM OF ALL FOUR CAMERAS SHOWING EQUAL INTENSITIES WHEN POINTED AT A BLANK WALL.	45
FIGURE 41: RAW FIREFLY MV IMAGES OF A MACBOOK PRO LAPTOP SCREEN DISPLAYING AN IMAGE. (TOP LEFT) CAMERA 0, (TOP RIGHT) CAMERA 1, (BOTTOM LEFT) CAMERA 2, (BOTTOM RIGHT) CAMERA 3.	45
FIGURE 42: RAW IMAGES OF A MACBOOK PRO LAPTOP SCREEN TAKEN AT A 45° ANGLE. (TOP LEFT) CAMERA 0, (TOP RIGHT) CAMERA 1, (BOTTOM LEFT) CAMERA 2, (BOTTOM RIGHT) CAMERA 3.	46
FIGURE 43: A SCREENSHOT OF THE SOFTWARE GUI FOR THE POLARIMETRIC CAMERA.	47
FIGURE 44: A TYPICAL SCENE LOOKING FROM RANDOLPH HALL TOWARDS WHITTEMORE HALL AT VIRGINIA TECH.	50
FIGURE 45: THE SUBTRACTED IMAGE FROM CAMERA 0 AND CAMERA 1.	50
FIGURE 46: A RAW IMAGE PLACED INTO A LARGER IMAGE ARRAY.	50
FIGURE 47: AN IMAGE SHOWING THE TEST SETUP USED TO DETERMINE PIXEL SHIFT AT VARIOUS DISTANCES.	51
FIGURE 48: THE CALIBRATION GRID AS OBSERVED BY CAMERA 0 (LEFT) AND CAMERA 1 (RIGHT) AT A DISTANCE OF FIVE FEET.	51
FIGURE 49: THE SUBTRACTED IMAGE FROM CAMERA 0 AND CAMERA 1 OF THE GRID SHOWN IN FIGURE 48.	52
FIGURE 50: A GRAPH SHOWING THE RELATIONSHIP BETWEEN THE AMOUNT OF PIXELS CAMERAS 1-3 HAD TO BE SHIFTED TO ALIGN WITH CAMERA 0.	52
FIGURE 51: A SET OF RAW IMAGES OF A CFL LIGHT BULB TAKEN WITH THE PROTOTYPE POLARIMETRIC CAMERA.	53
FIGURE 52: A SET OF IMAGES TAKEN LOOKING FROM INSIDE RANDOLPH HALL TOWARDS WHITTEMORE HALL AT VIRGINIA TECH.	54
FIGURE 53: A TOP-DOWN VIEW OF THE BEAM SPLITTERS, LABELED 1 THROUGH 3.	55
FIGURE 54: THE TRANSMISSION GRAPH FOR THE THORLABS 50:50 NON-POLARIZING CUBE BEAM SPLITTER [42].	55
FIGURE 55: THE SPECTRAL RESPONSE OF THE WHITE LED AIMED INTO THE BEAM SPLITTER ASSEMBLY.	56
FIGURE 56: THE SPECTRAL RESPONSE OF A WHITE LED OBSERVED FROM THE POINT OF VIEW OF CAMERA 0 (LEFT) AND CAMERA 3 (RIGHT).	56
FIGURE 57: A SET OF RAW IMAGES OF A -45° POLARIZED LCD MONITOR. (TOP LEFT) CAMERA 0, (TOP RIGHT) CAMERA 1, (BOTTOM LEFT) CAMERA 2, (BOTTOM RIGHT) CAMERA 3.	57
FIGURE 58: STOKES IMAGES OF A -45° POLARIZED LCD MONITOR. THE IMAGES ARE SHOWN ON A SCALE OF 0 TO 1, TRUNCATING ANY NEGATIVE VALUES FOR VISUALIZATION PURPOSES. (TOP LEFT) S_0 , (TOP RIGHT) S_1 , (BOTTOM LEFT) S_2 , (BOTTOM RIGHT) S_3 .	58
FIGURE 59: A SET OF RAW IMAGES OF A -45° POLARIZED LCD MONITOR ON THE LEFT AND A CINDERBLOCK WALL ON THE RIGHT. (TOP LEFT) CAMERA 0, (TOP RIGHT) CAMERA 1, (BOTTOM LEFT) CAMERA 2, (BOTTOM RIGHT) CAMERA 3.	59
FIGURE 60: STOKES IMAGES OF A -45° POLARIZED LCD MONITOR AND CINDERBLOCK WALL. (TOP LEFT) S_0 , (TOP RIGHT) S_1 , (BOTTOM LEFT) S_2 , (BOTTOM RIGHT) S_3 .	59

FIGURE 61: THE DEGREE OF POLARIZATION IMAGE OF AN LCD MONITOR ON THE LEFT AND A CINDERBLOCK WALL ON THE RIGHT.....	60
FIGURE 62: AN IMAGE OF THE TYPE OF MUDDY PATH ON THE DRILLFIELD AT VIRGINIA TECH AFTER HEAVY RAIN.	61
FIGURE 63: THE S0 STOKES IMAGE (LEFT), DEGREE OF POLARIZATION (CENTER), AND POLARIZATION CONTRAST (RIGHT) OF THE MUDDY PATH.	61
FIGURE 64: S0 IMAGE (LEFT), DEGREE OF POLARIZATION (CENTER), AND POLARIZATION CONTRAST (RIGHT) OF STANDING WATER ON THE DRILLFIELD.	61
FIGURE 65: S0 IMAGE (LEFT), DEGREE OF POLARIZATION (CENTER), AND POLARIZATION CONTRAST (RIGHT) OF STANDING WATER ON THE DRILLFIELD.	62
FIGURE 66: S0 IMAGE (LEFT), DEGREE OF POLARIZATION (CENTER), AND POLARIZATION CONTRAST (RIGHT) OF STANDING WATER ON THE DRILLFIELD.	62
FIGURE 67: STANDING WATER AND MUD ON THE DRILLFIELD WHEN VIEWED FROM ABOVE.	62
FIGURE 68: THE FOUR IMAGES FROM EACH CAMERA LOOKING DOWN STROUBLES CREEK.(TOP LEFT) HORIZONTALLY FILTER (TOP RIGHT) VERTICALLY FILTER (BOTTOM LEFT) 45° FILTER. (BOTTOM RIGHT) -45° FILTER.	63
FIGURE 69: THE STOKES VECTORS CALCULATED FROM THE IMAGES IN FIGURE 68. FROM THE TOP LEFT GOING CLOCKWISE: S0, S1, S3, S2.....	64
FIGURE 70: THE DEGREE OF POLARIZATION CALCULATED FROM THE STOKES IMAGES IN FIGURE 69.	64
FIGURE 71: THE DEGREE OF POLARIZATION IMAGE CALCULATED WITHOUT THE S3 TERM.	65
FIGURE 72: A POLARIZATION CONTRAST IMAGE OF STROUBLES CREEK.	65
FIGURE 73: THE DEGREE OF POLARIZATION OF STROUBLES CREEK WITH THE SUN BEHIND CLOUDS (LEFT) AND THE SUN VISIBLE (RIGHT).	66
FIGURE 74: THE POLARIMETRIC CAMERA AT STROUBLES CREEK.	66
FIGURE 75: STROUBLES CREEK VIEWED PERPENDICULAR TO THE SUN, LOCATED TO THE LEFT FROM THIS POINT OF VIEW.	67
FIGURE 76: THE DEGREE OF POLARIZATION OF STROUBLES CREEK VIEWED PERPENDICULAR TO THE SUN.	67
FIGURE 77: A PAIR OF IMAGE OF WEST CAMPUS DRIVE AT VIRGINIA TECH. (LEFT) HORIZONTAL POLARIZED IMAGE. (RIGHT) VERTICALLY POLARIZED IMAGE.	68
FIGURE 78: A POLARIZATION CONTRAST IMAGE OF WEST CAMPUS DRIVE ON VIRGINIA TECH’S CAMPUS.	68
FIGURE 79: A SET OF POLARIZED IMAGES OF A TYPICAL VAN. (TOP LEFT) HORIZONTALLY POLARIZED. (TOP RIGHT) VERTICALLY POLARIZED. (BOTTOM LEFT) 45° POLARIZATION. (BOTTOM RIGHT) -45° POLARIZATION.....	69
FIGURE 80: THE STOKES PARAMETERS CALCULATED FROM THE IMAGES IN FIGURE 79. (TOP LEFT) S0. (TOP RIGHT) S1. (BOTTOM LEFT) S2. (BOTTOM RIGHT) S3.....	69
FIGURE 81: RAW IMAGES OF CARS ON WEST CAMPUS DRIVE AT VIRGINIA TECH.	70
FIGURE 82: THE POLARIZATION CONTRAST IMAGE OF THE SAME SCENE SHOWN IN FIGURE 81.....	70
FIGURE 83: A CAD SCREENSHOT OF THE BEAM SPLITTER PLATFORM AND ASSEMBLY.	75
FIGURE 84: A SCREENSHOT OF THE CONTROL SOFTWARE WITH IMAGE ALIGNMENT PANELS HIGHLIGHTED IN BLUE AND ALIGNMENT IMAGE TABS HIGHLIGHTED IN RED.	85
FIGURE 85: A SCREENSHOT OF THE CONTROL SOFTWARE WITH CAMERA EXPOSURE CONTROLS HIGHLIGHTED BY THE BLUE BOX.	86

List of Tables

Table Heading	Page Number
The generic Mueller matrices for the polarimetric camera.	28
The polarizer-camera pairs in the polarimetric camera.	44
Mean intensities of each camera of the scene in Figure 51.	54
The mean intensities of the images in Figure 52.	54
List of components in the polarimetric camera and their part cost	74

1 Background and Motivation

1.0 Thesis Summary

This thesis will describe the process in developing, constructing, and testing a low-cost polarimetric camera. Development a low-cost polarimetric camera began after attempts to acquire a commercial polarimetric camera fell flat due to the high cost of entry, with cameras starting at over \$25,000 each [1]. A polarimetric camera is a sensor that is sensitive to the angle of light that is being reflecting away from objects. By comparing the polarization of different areas of an image, objects of interest can be segmented or identified, which could provide new and useful information for an unmanned ground vehicle to make better decisions in its navigation and obstacle avoidance. For example, light reflected off of water and mud is commonly assumed to be horizontally polarized. If one looks at the water or mud with cameras equipped with linear polarizers oriented horizontally and vertically, a difference in the intensity of the water's surface can be observed.

The low-cost design used for the camera in this thesis was a multi-camera system that used a set of beam splitters to optically align each camera. A beam splitter is an optical device that splits incoming light into two separate beams. With three beam splitters arranged in the right order and orientation with respect to each other, four cameras can look and see exactly the same scene, which is required to directly compare the images to each other. In front of each camera is a linear polarizing filter each set at a unique angle. By combining the slightly different information from each camera, the system can identify the state of polarization of the scene. For an unmanned ground vehicle, a polarimetric camera could provide new information at relatively low cost compared to other sensors such as LIDARs.

The polarimetric camera that was constructed was able to provide good results showing the polarization states of various scenes and objects, including standing and running water, muddy paths, and vehicles. Extra steps had to be taken in software to perfect the image alignment between cameras. Also, the beam splitter assembly induced an intensity difference between cameras that had to be manually corrected before taking data. However, once the alignment and intensity differences were corrected before a data capture, the camera was able to capture and calculate the linear Stokes parameters needed to characterize the state of polarization in real time and show them to a user.

1.1 UGVs at Virginia Tech

This section will briefly review the evolution of unmanned ground vehicles developed at Virginia Tech, with a focus on the perception systems used. Full sized ground vehicle development at Virginia Tech started in earnest in the DARPA Grand and Urban Challenges and continued with the GUSS UGV program for the US Marine Corps. The majority of the sensors used on these vehicles were single and multi-plane LIDARs that provided both object detection and some identification. Work on advanced perception with cameras began towards the end of the GUSS program by adding multispectral machine vision cameras to classify vegetation in a scene. Further work has been done on identifying people and water. The polarimetric camera would be an additional camera sensor that could be added to the advanced perception camera

suite. Polarization data is currently not collected by any of the sensors on any of the unmanned ground vehicles developed at Virginia Tech to date. The addition of a polarimetric camera would enhance the data available to a UGV's autonomy software and enable it to make more informed decisions.

1.1.1 2004 DARPA Grand Challenge

The Mechatronics Lab at Virginia Tech has a several year history of developing unmanned ground vehicles, starting with the 2004 DARPA Grand Challenge. The 2004 DARPA Grand Challenge was the first robotics competition issued by DARPA to jumpstart unmanned ground vehicle development by offering cash prizes to teams. Teams had to complete a 150 mile road course in the Mojave Desert in California without any human intervention.

Virginia Tech entered Cliff, a converted four-wheel drive Club Car utility vehicle, into the 2004 Grand challenge [2]. Cliff carried radar, LIDAR, and visual cameras to detect obstacles in its path. The LIDAR sensors were placed so when combined, they could measure the ground at different levels. The absolute difference between measurements taken from the roll cage versus measurements taken from the front bumper allows the vehicle to decide if an object is present in front of the vehicle or if the ground slope was gradually changing. The combined data was used to create a local terrain map and determine the best immediate path for Cliff to take. The radar and camera sensors ended up not being used in the competition run [2].

Virginia Tech was one of 15 teams that qualified for the final event. However, no teams were able to travel more than 7 miles due to mechanical problems or navigation error. Cliff traveled approximately 30 meters before the computers applied the brakes and did not let it advance any further.

1.1.2 2005 DARPA Grand Challenge

In 2005, DARPA hosted another Grand Challenge event to see if teams could improve upon the results from 2004. Virginia Tech entered two vehicles, Rocky and Cliff, shown in Figure 1. They were designed independently with different navigation schemes to see which vehicle performed better and to compare the two directly [3].




Figure 1: Virginia Tech's two entries into the 2005 DARPA Grand Challenge, Cliff (left) and Rocky (right) [4].

Rocky and Cliff had the same base vehicle platform and sensor suite. Two downward facing single plane Sick LIDARs and a third mounted horizontally on the bumper were used for local

terrain mapping and obstacle detection. A Point Grey Bumblebee stereovision camera is used to observe the area directly in front of the vehicle. Using the color data from one of the visual cameras, a color value for a road was logged once the vehicle determined where the road was. It did this by checking feature such as edge continuity, vertical narrowing, and relative size in the image. The pixel location of the road in the image was then converted to coordinates relative to the vehicle frame to be used in path planning.

Rocky and Cliff traveled 39 and 44 miles on the course, respectively. Both vehicles failed due to mechanical, rather than navigation problems. Rocky's onboard generator shut down due to a false low oil reading and Cliff's engine stalled when it briefly idled.

1.1.3 2007 DARPA Urban Challenge

The 2007 Urban Challenge was the third unmanned ground vehicle competition put on by DARPA to further advance the field of robotics and autonomous vehicles. The main difference between the 2007 challenge and the 2004 and 2005 challenges was the environment. The 2007 Urban Challenge was set in a model town part of the Victorville, CA airport. The course was 60 miles long, with a target completion time of less than 6 hours. Teams were given GPS locations of checkpoints that the vehicles had to drive through 24 hours beforehand.

The course had static obstacles such as buildings, signs and utility poles, as well as moving obstacles, such as other autonomous and human controlled vehicles. The autonomous vehicles running the course had to follow all California driving laws and could only pause for up to 10 seconds at a time. Since the vehicles did not operate isolated from one another, they had to make real time decisions based on what their perception systems were telling them. For example, the vehicles had to correctly navigate an intersection with multiple vehicles at a 4-way stop.

Team Victor Tango, a combination of engineers from Virginia Tech and TORC Technologies, was one of 11 teams that qualified for the national event in Victorville, CA. The vehicle entered was a 2005 Ford Escape Hybrid named "Odin" [5]. The Ford Escape platform, shown in Figure 2, was chosen because the stock steering, shifting, and throttle were already drive-by-wire, reducing the time needed to convert the vehicle into a fully drive-by-wire platform. In addition, DARPA required that the OEM safety systems could not be bypassed. Since the drive-by-wire systems did not require any modification, these requirements were easily met. The hybrid platform meant that the on-board electrical system was more than capable of handling the added load that sensors and computers would place on the vehicle without any additional modification necessary.



Figure 2: The Victor Tango Ford Escape hybrid, Odin, with sensors labeled [5].

Odin had several sensors to navigate through the course and to see its environment. Its main localization sensor was a NovAtel GPS/INS unit. As the vehicle moved, the NovAtel would update its position in the world reference frame. To detect obstacles, a set of LIDARs were used. The sensing areas around the vehicle are shown in Figure 3.




Figure 3: The sensor coverage around Odin [5].

Odin had two IBEO Alasca XT LIDARs that covered a 270° field of view in front of and to the side of the vehicle. Two SICK LMS LIDARs are positioned to look downwards from the roof rack to observe the areas immediately in front of and off the front corners of the vehicle. Two other SICK LMSs are positioned to look in the blind spots to the sides of the vehicle. For looking behind the vehicle when reversing and to observe any approaching vehicles, an IBEO Alasca A0 was used.

Two visible cameras were included on the roof rack for classifying objects, which the LIDAR-only systems could not do. Objects were split into two groups, ones that will not move and ones that are moving or could move. The visible cameras were deemed too computationally expensive

to use, especially at higher speeds. Thus, the IBEO LIDARs were the only sensors used to classify objects.

Odin did very well at the Urban Challenge, coming in third place overall with a time of 04:36:38, one of three teams to come in under the 6 hour goal set by DARPA and one of six teams to finish the course.

1.1.4 GUSS Program

The Ground Unmanned Support Surrogate (GUSS) program took the progress made on the DARPA Grand Challenge and Urban Challenge vehicles and then went to apply them to real world problems for a customer with specific needs. The GUSS project was a group of prototype UGVs developed for the United States Marine Corps under the supervision of NSWC Dahlgren for the purpose of assisting squad-level operations for deployed Marines. A Polaris ATV was chosen as a stand-in for an actual military vehicle. The GUSS vehicles, shown in Figure 4, would be used to determine how a UGV would most effectively be used by the military when out on a mission.



Figure 4: Three GUSS unmanned ground vehicles [6].

The GUSS vehicles operated in three main modes. First was a “follow me” mode, where the vehicle would closely follow a Marine on foot through the environment. The main motivation for this mode was to reduce the carrying load that each individual Marine normally carries, which could reach up to 100lbs of water, food, ammunition, and batteries [6]. By having a UGV transport these supplies with the squad, the Marine would expend less energy when out on a mission and be more capable and alert on long patrols and more mobile in dangerous situations.

The second and third was a target mode and a teleoperated mode. Target mode tells the vehicle to move to a certain spot and teleoperated mode allows a user to directly control the vehicle’s movements [7]. In these modes, a vehicle out on a mission with a squad could be sent to return back to a forward operating base (FOB) or another squad. A UGV that can autonomously go to a certain point or return to base could be used to shuttle supplies between separate squads without exposing a driver to danger, such as on the many IED attacks on vehicle convoys in the Iraq and Afghanistan conflicts. A UGV could also transport an injured Marine away from the front lines back towards the FOB for immediate treatment. Both situations currently require a Marine to drive the vehicle, whereas a UGV could keep that driver with the rest of the squad, reducing the amount of time spent exposure to danger during transport.

The GUSS vehicles used LIDAR to perceive their environment and detect obstacles, like the DARPA Grand and Urban Challenge vehicles. While this worked for the majority of environments and situations that the GUSS vehicles encountered, there were a few specific cases that caused unwanted behaviors. For example, tall grass growing in between vehicle ruts on a path caused the LIDAR systems to think that there was a pole sticking up out of the ground. While the vehicle could easily drive through the grass, it instead stopped and tried to find a way around. These situations where the LIDAR perception system breaks down prompted Virginia Tech and the Marine Corps to start research into additional sensors that could give the vehicle the extra information it needed to make better decisions. Section 1.2.2 describes the use of passive sensing to defeat these problems.

1.2 UGV Sensing Methods

There are two general methods to sense the environment around a vehicle: active and passive sensing. Active sensing relies on emitting energy from a sensor, then reading a response. Active sonar and radar are examples of active sensors that emit energy, sound and radio waves, in these cases. Passive sensing, such as cameras, does not require any energy emission. The polarimetric camera is a specific kind of passive sensor that not only measures the intensity of light of a scene but also can calculate the state of polarization of areas in the scene.

1.2.1 Active Sensing

For an unmanned ground vehicle, one of the primary active sensors used is LIDAR, which stands for Light Detection and Ranging. The LIDARs used on UGVs at Virginia Tech work on the time of flight principle. A laser beam is emitted from the sensor, which strikes an object and returns back. The time it takes between sending the laser and receiving it can be measured and converted into a distance using the speed of light. By sweeping a laser across a plane, LIDAR sensors can detect potential objects in front of it. A multi-plane LIDAR scan is shown in Figure 5.

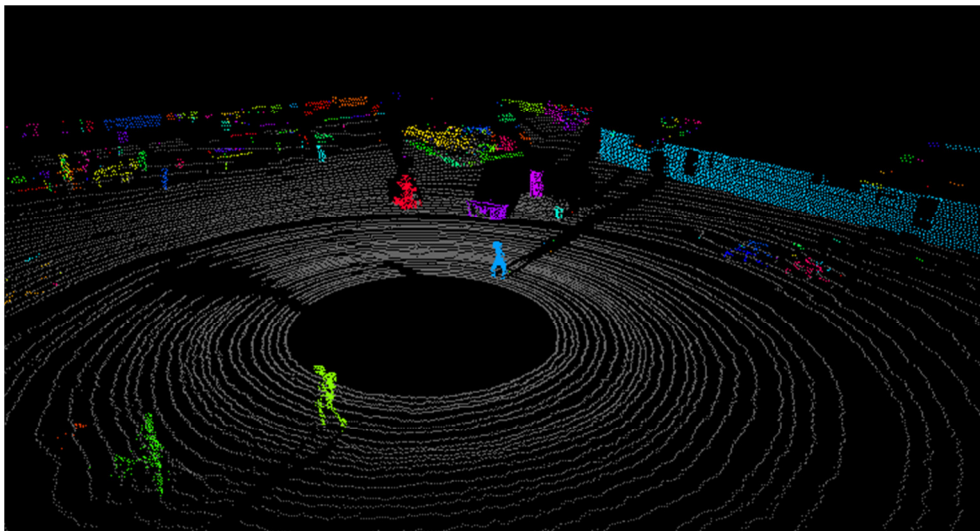


Figure 5: A Velodyne 64-E LIDAR scan with objects segmented as different colors [8].

In Figure 5, the grey lines each represent a scan plane. This particular LIDAR sensor sweeps 64 planes 360° and reads hundreds of thousands of points of data per second. Points that stand out from the rest of the scan have been segmented and colorized as objects of interest. By analyzing

the shapes that the segmented areas take, a computer system can potentially determine the difference between a wall, a person, a road, or other objects of interest. The active aspect of the technology means that the signal to noise ratio is very good, ensuring accurate measurements in the vast majority of environments and situations it may encounter.

While LIDAR remains the primary sensor on many unmanned ground vehicles, they have several drawbacks. LIDAR sensors only return distance information. Two objects of similar shape and at the same range will look similar to the LIDAR, regardless of what the actual makeup of the object is. For example, a human lying down on the ground can appear to have the same size and shape as a mid-sized log. To a UGV with only LIDAR sensors, it cannot tell the difference. LIDAR point cloud data is also sparse when compared to an image. Only areas that are struck by a laser are measured, leaving any area between points as unknown.

In addition, LIDAR sensors are currently still relatively expensive, when compared to vehicle cost. A single plane LIDAR that can withstand outdoor environments can cost several thousand dollars, with multi-plane LIDAR sensors quickly escalating into the tens of thousands of dollars. If UGVs are to become more commonplace, the cost of sensing has to decrease.

Finally, because LIDAR is an active sensor, it is susceptible to being detected by a third party. For a military vehicle, this is an unwanted behavior, as it increases the possibility of an enemy force detecting the vehicle at a longer distance than it would without a LIDAR. This could reduce the effectiveness of a military UGV and possibly place it in danger of being attacked and destroyed.

1.2.2 Passive Sensing

Traditional camera based perception systems measure the total intensity of light being reflected from a scene. Different objects reflect different wavelengths of light depending on their composition. Using this information, a computer system could identify what an object in an image is, or what it is made of. Compared to the cost of LIDAR sensors, many machine vision cameras are relatively inexpensive and potentially add a significant amount of information.

There is a wide range of wavelengths that camera systems can operate in. Most cameras operate in the visible wavelengths that humans see in, 400nm to 700nm. Also available are ultraviolet, near-infrared, medium wave infrared, and long wave infrared. These different wavelengths allow for different features to be detected from the same object and for object classification and identification to take place.

For example, a multispectral camera system can identify chlorophyll-rich vegetation with good accuracy. The relationship, called the Normalized Difference Vegetation Index (NDVI), uses the difference in intensity between near-infrared and red to determine if an object in a scene is likely vegetation. An example of NDVI segmentation can be seen in Figure 6.




Figure 6: A set of three scenes taken with visible (left) and near-infrared (center) sensors. The NDVI segmented image is on the right.

The calculations needed for NDVI are computationally inexpensive and potentially give a UGV information it does not have with a LIDAR-only system. For example, if a UGV had a LIDAR-only perception system, a tall blade of grass could potentially look like a guy wire or a pole, stopping it from proceeding further. If the UGV system had a multispectral camera as a part of its sensor suite, NDVI could be used to identify the grass and tell the UGV that it is safe to continue moving. A prototype NDVI system was applied on a GUSS vehicle, with promising results in defeating false positives from the original LIDAR-only perception system.

There are many more scenes and situations that a camera-based perception system can be used for, such as human detection and water detection. Machine learning algorithms can combine information from many different sources, such as multispectral cameras and LIDAR sensors, to make decisions based on statistics from previously observed and classified data [9] [10].

In addition to performing object identification and classification, vision data can be used to identify distance to an object. Ranging is achieved using cameras by implementing stereo vision. Stereo vision is the same mechanism allows humans to see depth with two eyes. A stereo vision depth map can be seen in Figure 7. By comparing the difference in images from two offset cameras, a depth can be calculated and correlated with a distance. This is a direct competitor to a LIDAR sensor that also returns distance information.




Figure 7: An outdoor scene taken with a near-infrared camera (left) and the corresponding depth map (right).

The stereo depth map on the right in Figure 7 colors pixels based on how near or far away the area of the scene is. Red is classified as close to the camera and blue is far, with varying shades in between. Stereo depth information is much denser than LIDAR, but is more susceptible to noise because it is a passive system and relies on many variables that can affect accuracy.

However, sensor cost for a stereo system can be orders of magnitude lower than a LIDAR and work across many different wavelengths. It is not inconceivable for a camera-only perception system that returns depth in addition to color and intensity to succeed in place of a LIDAR-based perception system. Currently, many UGVs are combining LIDAR data with vision data to get the benefits of both modalities.

1.3 Why Polarized Light?

The advanced perception team initially experimented with polarizing filters as a method to reduce glare and reflections from water in visible and infrared images. After attending the US Army Hyperspectral and Polarimetric Workshop in Huntsville, AL in July 2012, the advanced perception team began looking at quantifying the polarization of light directly to be used as a new source of information for a UGV.

Efforts to obtain a commercial polarimetric camera were quickly halted due to the high cost of purchasing a camera, with prices starting in the \$25,000 range and escalating quickly. Companies that make polarimetric cameras were also hesitant to loan out cameras to evaluate their performance without purchasing. The idea to construct a polarimetric linear Stokes camera soon followed.

Field tests using a single JAI AD-080GE multispectral camera with a standard linear polarizer showed potential in segmenting parts of an image such as reflections of standing water. Figure 8 and Figure 9 show the results of one of these early tests.



Figure 8: A visible image of the Duck Pond at Virginia Tech.



Figure 9: A thresholded degree of polarization of the same scene from Figure 8.

The polarized images were taken by manually rotating a linear polarizer to specific angles before saving each image. They were processed afterwards in MATLAB to calculate the degree of polarization over the entire image using the linear Stokes parameters. The processed images clearly showed that the surface of the water was differently polarized than much of the surrounding environment. It also showed that the sky had different polarization characteristics than the vegetation and the ground. The thresholded image also showed that polarization was not a function of what the surface of the water was reflecting from the point of view of the camera. The water surface is reflecting both sky and trees in the visible image, but both are relatively highly polarized in the processed image.

1.4 Thesis Content

Chapter 1 of this thesis described the historical background and motivation for why the use of polarized light could be useful for an unmanned ground vehicle. It reviewed the history of unmanned ground vehicles at Virginia Tech. It also went into detail what the vehicles had to sense in their environments and the sensor packages used to do so.

Chapter 2 will go over the literature review for different areas of interest for the polarimetric camera. It will review the history of how the polarization of light was first discovered and used.

Chapter 2 will also go over how polarization is used by insects to both navigate in their environment as well as distinguish water from land and mirages.

Chapter 3 will review the methods used to design and assemble the polarimetric camera itself. It will first go over the fundamental science behind polarized light and the optical hardware used. Then it will go over the detailed design and assembly of the polarimetric camera used in this thesis.

Chapter 4 will review the experimental data taken with the polarimetric camera. It will first go over tests done indoors in a controlled environment to characterize the state of polarization, check intensity difference between cameras, and check image alignment between cameras. It will then take the camera outdoors to view scenes that include running water, standing water, mud, and automobiles.

Chapter 5 will go over the results from the data collected in Chapter 4 and discuss areas where the camera worked and did not work as designed. It will also discuss how the polarimetric camera could be used on a future UGV as an additional sensor.

Chapter 6 will conclude the thesis and summarize the results and future work for the prototype polarimetric camera system

2 Literature Review

2.0 Chapter Summary

Chapter 2 of this thesis will review the history of polarized light and how it was first discovered. The chapter will then summarize how insects use polarized light to navigate and to identify water. It will then summarize the different methods of observing the state of polarization of incoming light using a camera in both research and commercially available systems. Finally, the chapter will discuss image processing that uses polarization data to classify areas of an image as water or mud, as well as filter an image for haze.

2.1 Discovery of Polarization

The first rumored use of the polarization of light was in the 13th and 14th centuries in Iceland. A medieval text, *The story of Rauð and his sons*, makes reference to a “sun-stone” that was used to determine the position of the sun on an overcast day [11]. However, there is little evidence supporting the theory that Vikings used this method as a form of localization. In 2012, researchers at Eötvös University in Budapest determined that the use of polarization as a locating mechanism for the sun was theoretically possible in the Arctic regions that Vikings would have sailed in [12].

The first documented discovery of the polarization of light occurred in 1690, when Christian Huygens demonstrated that the refracted light from a crystal of calc-spar (or Iceland Spar), a transparent form of calcite, could be extinguished by passing it through a second crystal rotated about the direction of the incoming wave of light [13]. The first demonstrated method of polarizing light by reflection was discovered by Étienne-Louis Malus in 1808. Malus reflected a beam of unpolarized light off a glass surface at an angle. The light was then reflected again by a second glass surface parallel to the first. If the second glass surface was then rotated about the axis of the incoming light, the intensity of the second reflected beam was observed to decrease until it reached zero when the second plate rotated to 90° [13], [14]. Both methods linearly polarized the light when it passed through the first crystal or reflected off the first glass surface. When the light interacted with the second surface, the 90° difference in orientation blocked the light from passing through.

2.2 Polarization in nature

Several species of insects have been observed to be sensitive to polarization. The insects use their ability to observe polarized light to both navigate and distinguish different areas of their environment from one another.

2.2.1 Honeybee orientation with skylight polarization

A well-known example of the use of polarization by insects is the honeybee *Apis mellifera*. In 1949, Karl von Frisch observed that the dances of bees became disoriented in total darkness, but corrected themselves and reoriented when a light source such as the sun, blue sky, or a lamp became visible. The bees use the dances to communicate where sources of food are located with respect to the hive. When the bees were dancing on a vertical surface in total darkness, the bees oriented with respect to gravity. However, with about a 10 to 15° patch of blue sky visible, the bees orient themselves toward the direction of the sun, even if the sun was not directly visible

[15]. Frisch theorized that sky-light polarization could be the means by which the bees oriented themselves. To test the hypothesis, he linearly polarized the light in a controlled setting on bees dancing on a horizontal surface. When the polarizer was turned, the direction of the bees' dance changed along with it. Further experiments determined that the bees were sensitive specifically to polarization in the UV and blue bands. The sensitivity in the UV and blue bands allows the bees to orient themselves in an overcast sky, as the sun is still visible in the UV band through cloud cover.

2.2.2 Aquatic insects and water detection

There is ample evidence to suggest that aquatic insects use polarization to distinguish water from the rest of the environment. A study done on mirages found that there is no contrast in the degree and angle of polarization between the sky and a mirage but a large contrast between the water surface and the sky [16]. Bodies of water tend to reflect and horizontally polarize light, while mirages gradually bend the light back upwards without changing its polarization. Polarization-sensitive insects then use this difference in observed polarization to distinguish bodies of water from a mirage [15]. To a non-polarization sensitive observer, such as a human, a mirage and a body of water could potentially look the same.

A study on insects' attraction to bodies of water found that insects perceive the polarization of light reflecting off water at a wide range of wavelengths from the UV to the red [17]. To determine what wavelengths different insect species used to identify water, the study determined how the degree of linear polarization varied with wavelength as shown in Figure 10.




Figure 10: Degree of linear polarization P of light reflected at the Brewster Angle from a pond with clear water and dark bottom (dashed line) and a eutrophic pond with green-brown appearance (solid line) [17].

The study found that the degree of linear polarization is higher at shorter wavelengths than at longer wavelengths.

2.2.3 Human Polarization Sensitivity

Human sensitivity to polarization was first discovered by Wilhelm Karl von Haidinger in 1844.[15] The effect, called Haidinger Brushes, can be seen if one gazes at strongly linearly polarized light emitting from an unstructured area. For example, looking at a uniform background through a linear polarizer or at a blank white area on an LCD screen can produce the

effect. Haidinger Brushes take the form of a faint bow-tie or propeller shaped pattern aligned with the polarization.

2.3 Polarization Cameras

This section will discuss commercially available polarization cameras as well as ones developed for research. Polarization cameras can use different methods to observe a scene at different polarizations. These methods include division of time, beam splitters with linear polarizing filters, and beam splitters with wave plates and retarders.

2.3.1 Division of Time polarization camera

There are a handful of commercially available linear Stokes polarization cameras available on the market today. Bossa Nova Tech sells a compact linear Stokes polarization camera for use in research. It is a division of time system, meaning that one sensor is used with rotating or changing polarizing filters placed in front of it. The Bossa Nova SALSA camera uses a set of two polarization rotators. The first set has a quarter wave plate and a programmable quarter wave plate and the second has a programmable half wave plate and a polarizer. The polarization states of the camera are shown in Figure 11.

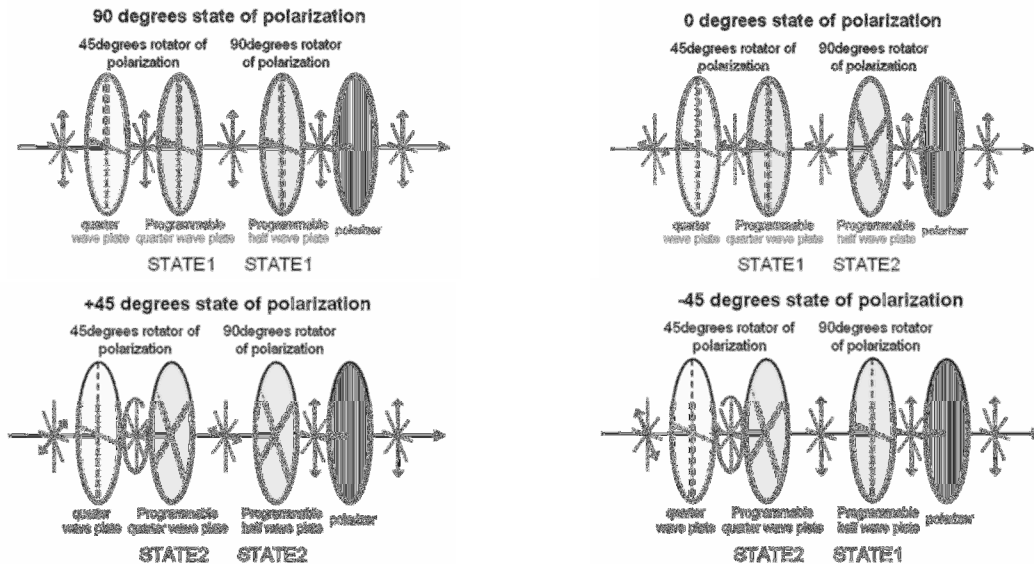


Figure 11: Polarization module principle of the SALSA camera [18].

By changing the orientation of the programmable quarter wave and half wave plate, polarization states of -45° , 0° , 45° , and 90° can be achieved with a small time delay between frames [18]. This method is faster than using a physically rotating polarizing filter to achieve the states and has the added benefit of only requiring one sensor to record data. With only one sensor, a smaller overall unit size can be achieved. However, only one sensor means that multiple frames are needed to fully characterize the polarization, which could potentially pose problems with a moving UGV. According to emails from Bossa Nova Tech, the SALSA camera can cost up to \$25,000[1].

2.3.2 Polarization camera using a beam splitter

The other method of sampling multiple polarizations involves the use of a beam splitter and independent imaging sensors. Using a beam splitter, a sensor can sample the same area of an image simultaneously through different filters. Types of beam splitters available include cube

beam splitters, pellicle beam splitters, and plate beam splitters. Cube beam splitters often have a limited field of view and pellicle beam splitters are sensitive to vibration, which limits their use on a moving vehicle, leaving plate beam splitters as an ideal choice [19]. The basic layout of the beam splitter polarization camera can be seen in Figure 12.




Figure 12: Top-view diagram of a two-CCD polarization camera using a polarizing beam splitter [19].

Incoming light is reflected and transmitted with a plate beam splitter oriented at 45° . Two cameras are placed in the path of the reflected and transmitted light, which optically aligns them and allows for pixel to pixel correspondence between both cameras.

This specific setup used a polarizing plate beam splitter and an electronically controlled crystal that could shift the phase of the light by 90° . Flipping the phase shift back and forth as both cameras took a frame allowed for the full polarization state to be calculated after just two frames. Assuming that the polarization parallel to the page is P and the polarization vertical to the page is S, then the intensity of the polarized transmitted and reflected light can be represented as

$$S = \frac{I_{transmitted}(1 - a) - aI_{reflected}}{b - a} \quad (2.1)$$

$$P = \frac{I_{transmitted}(1 - b) - bI_{reflected}}{a - b} \quad (2.2)$$

where $a^2 + b^2 = 1$, $a \geq 0$, $b \geq 0$, $a \neq b$. Coefficients a and b are dependent on the coating of the beam splitter used [19].

A commercially available polarization camera that uses a beam splitter and linear polarizing filters is available from FluxData [20]. The FD-1665 is a three-sensor multispectral camera that can have polarizing filters inserted before each sensor, as shown in Figure 13. It uses a 3-way prism beam splitter to send the light to each filter and sensor.




Figure 13: Schematic view of the Fluxdata 3-CCD camera showing the 3-way beam splitter and polarizer configuration [21].

The FluxData FD-1665 can be configured as a monochrome or color camera, with polarizers typically oriented at 0°, 45°, and 90° or 0°, 60°, 120°. The Stokes parameters can be easily calculated with the intensities from each sensor. However, according to emails from FluxData, an FD-1665 configured as a polarization camera can cost between \$25,000 and \$50,000 to purchase, depending on what sensors are used [22].

2.3.3 Polarization camera using a beam splitter and waveplates

Another method of observing polarized light is to use beam splitters in combination with wave plates to observe different states of polarization. A beam splitter polarization camera using this method was constructed to observe the slope of waves in the ocean [23]. Four cameras were aligned around a set of three partially polarizing beam splitters, as shown in Figure 14.




Figure 14: Schematic of a beam splitter polarization camera that uses wave-plates [23].

The first beam splitter transmitted 80% p-polarized light and 20% s-polarized light and reflected 20% p-polarized light and 80% s-polarized light. The transmitted light was passed through a half-wave plate oriented at 22.5° , then through a second polarizing beam splitter that transmitted the p-polarization onto one sensor and reflected the s-polarization onto a second sensor. The reflected light from the first beam splitter was passed through a quarter-wave plate oriented at 45° , then passed through a second polarizing beam splitter that transmitted the p-polarized component onto the third sensor and reflected the s-polarized component onto the fourth sensor.

For this system, S_0 was defined as the sum of intensities across all four sensors. S_1 is the difference of intensities between the sum of paths 1 and 2 and the sum of paths 3 and 4. S_2 is the difference of intensities between path 1 and path 2. S_3 is the intensity of path 4 minus the intensity of path 3. These were the four Stokes parameters needed to determine the polarization characteristics of the incoming light. Once built, the camera was able to run at 60 frames per second, which was necessary to be able to characterize quickly changing waves in the ocean.

2.4 Scene Classification and Image Processing

2.4.1 Mud detection using polarization

The ability to detect mud would be highly desirable for an unmanned ground vehicle. Rankin and Matthies used polarization contrast as a method of detecting mud [24]. Polarization contrast only requires two images of a scene taken with a polarization offset of 90° . The method assumes that one of the two polarization images taken is aligned with the orientation of the incoming linearly polarized light. If the incoming linearly polarized light is not aligned with one of the two polarizing filters, then polarization contrast provides a lower bound to the degree of linear polarization instead.

Since water and mud were assumed to horizontally polarize the light reflecting off it, the linear polarization angles chosen to be measured were 0° and 90° . To calculate the polarization contrast at each pixel, the absolute value of the difference of the two images were divided by the sum of the two images.

$$PC = \frac{|I_0 - I_{90}|}{I_0 + I_{90}} \quad (2.3)$$

In the first experiment, a large patch of mud was observed over the course of seven hours. Matthies found that the polarization contrast of the mud did not change over the course of the day and that the polarization contrast of the mud was significantly higher than the surrounding dry soil, as shown in Figure 15.




Figure 15: A set of polarization contrast and segmented images of a patch of mud taken seven hours apart [24]

The same research also looked at the degree of linear polarization of mud from various orientations. It found that the degree of linear polarization remained high regardless of the direction the mud was observed from. However, standing water had a much lower degree of linear polarization than the wet soil, since the degree of linear polarization is dependent on sky conditions and the orientation of the sensor relative to the sun. Finally, mud that was in shadow was hard to distinguish from the rest of the soil using the degree of linear polarization [24].

2.4.2 Water detection using passive sensing

Passive water detection using cameras is a subject of intense study in the field of unmanned ground vehicles. Research has been done on using color, stereo vision, short-wave infrared, thermal infrared, and polarization as potential methods of detection [25].

Color imagery is often used due to the relatively low cost of high quality machine vision cameras. Color imagery gives an algorithm data such as brightness and saturation. On average, the mean brightness of the sky tends to be at near saturation of the pixel values and that the sky is two and a half times brighter than the terrain. The average brightness of water is mid-way between the sky and terrain [26]. A machine learning algorithm can use that relationship to classify pixels in images with some accuracy. According to Matthies, incorrect classifications tended to occur when an image of vegetation was reflected in still water, which is a problem that has also occurred in the NDVI images described in Section 1.2.2.

Stereo vision could also be used to identify bodies of water. Stereo vision is the method of using two cameras to provide depth information by measuring the disparity between images of the same object or feature in a scene, similar to how a human sees depth with two eyes. In a scene that has water, the depth map could show reflections from the water's surface as depths that are

farther away than the surrounding terrain, or below the ground plane [26]. This information can be combined with a method of detecting water using sky reflections if viewed at large incidence angles to detect bodies of water with good accuracy [27].

Outside of the visual spectrum, short-wave (SWIR) and long wave infrared (LWIR) imaging are two non-visible methods of segmenting water. Water heavily absorbs SWIR at around 1450nm, making it appear very dark or black when viewed through a SWIR camera. Bodies of water tend to be warmer than the surrounding land at night and cooler during the day, which is a characteristic easily measured with a LWIR camera [26]. These known characteristics can be used to identify potential areas of water with simple intensity-based segmentation techniques.

According to David West, the polarization of snow has shown to not be as useful of a descriptor compared to intensity due to the highly varied surface and the Umov Effect. The Umov effect states that bright surfaces tend to have unpolarized reflections while dark ones have highly polarized ones [28]. The same paper showed a small correlation to changes in the degree of linear polarization between disturbed snow, such as footsteps or tire tracks, and undisturbed snow.

2.4.3 Haze Reduction

Polarization can also be used to reduce haze present in an image. Schechner used polarization contrast to reduce haze and improve image quality beyond what simple optical filtering can accomplish [29]. The process assumes that the polarized light reaching the sensor is made up from two sources, the direct transmission and the airlight. The direct transmission is the scattered signal reaching the camera from the scene objects and the airlight is the ambient illumination from the sun scattered by the air. Both sources of light have increased scattering as distance increases.

Schechner assumes that the direct transmission portion of the light is unpolarized. Thus, variations of intensity as polarizer filter angle changes are assumed to be due to airlight. The irradiance due to airlight is measured as a function of polarization angle and is used to determine what the best case and worse case angles are. The airlight component of the irradiance is considered to be at its highest, or worst, when the polarizing filter is horizontal and is the lowest, or best, when vertical. The worse state case occurs because the irradiance due to the airlight is enhanced compared to the irradiance from direct transmission, which is assumed to be constant. Using the difference between the two, the irradiance due to the airlight can be estimated. Using the two raw images, the dehazed image can be calculated using Equation (2.4). An example is shown in Figure 16.

$$\hat{I}^{object} = \frac{\hat{I}^{total} - \hat{A}}{\hat{t}} \quad (2.4)$$

\hat{I}^{object} is the dehazed image, \hat{A} is the estimated airlight, and \hat{t} is the estimated transmittance [29].

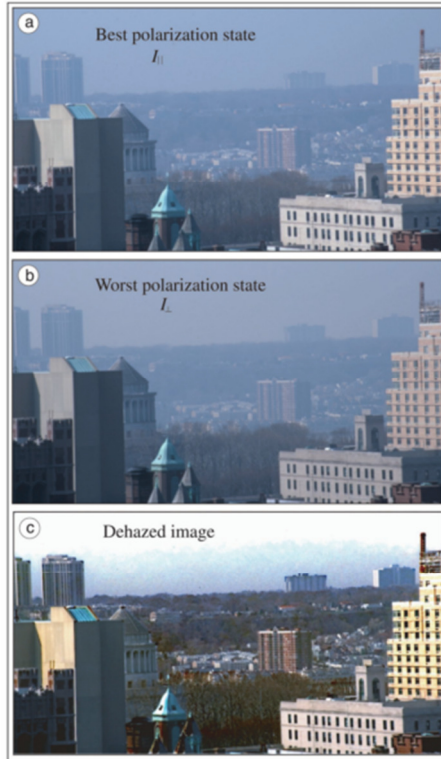


Figure 16: (a) A raw image taken at one polarization state. (b) A second raw image taken at a different polarization state. (c) The final processed image with haze reduced [29].

2.5 Thesis Relevance

The literature review in Chapter 2 identifies the polarization state of light as a potentially useful source of information. There are examples of polarization being used by insects to locate the sun and navigate on cloudy days and to identify water. This example shows the potential to do similar sensing on an unmanned ground vehicle. A UGV that can use polarization to sense where it is or perform obstacle navigation will have advantages over a UGV that does not have sensors that can detect this information. A UGV could add a polarimetric camera at relatively low cost and greatly enhance its ability to perceive its environment in some of the more difficult navigation cases.

For example, a UGV that approaches a large patch of mud or water with only LIDAR or traditional camera sensors may not be able to differentiate between the mud and the surrounding ground. While a human driver can tell where potential difficult spots are and avoid them, a computer will have a much harder time differentiating with a traditional sensor suite and could potentially get the vehicle stuck. With a polarimetric camera capable of identifying these previously unidentifiable areas, an autonomous UGV would have its decision making capabilities improved.

The literature review included commercially available cameras that can accomplish everything a UGV would need in a polarimetric camera, except at a cost equal to a full size vehicle. This thesis will determine the feasibility of developing a low cost polarimetric camera and what capabilities it would have.

3 Methods of Camera Design and Assembly

3.0 Introduction

Chapter 3 is split into two main sections. The first section will discuss the fundamental science behind light, including a summary of the electromagnetic wave model and how the polarization of light is a special case. It will also discuss how the polarization of light is mathematically described and how the polarization state of a scene can be calculated using the Stokes parameters. The second section of chapter 3 will discuss the detailed design of the polarimetric camera including its design requirements, selection and layout of components, and the issues that came up during the assembly of the camera.

This thesis aims to develop a polarimetric camera at relatively low cost. This thesis will evaluate the design process and potential problem areas in going with an inexpensive route. By determining what data could be collected with a low cost system, a future UGV design that wants to use polarimetric data can make a more informed decision about purchasing a commercial system or developing one in-house.

3.1 Fundamental Science and Optics Hardware

This section will discuss the fundamental science behind polarized light as well as discuss the optical hardware used in the polarimetric camera. The fundamental science of polarized light will discuss the electromagnetic wave model and how the orientation of the electric field relates to the polarization state of the light. It will also discuss mathematical methods for describing the polarization state and how they could be used to process images taken with the polarimetric camera.

3.1.1 Electromagnetic Waves and Polarized Light

Light is a transverse electromagnetic wave. Transverse waves have oscillations that are perpendicular to the direction of the wave's travel. Electromagnetic waves are comprised of two components, an electric field, \vec{E} , and a magnetic field, \vec{B} , shown in Figure 17.




Figure 17: An illustration of an electromagnetic wave [30].

The two fields are perpendicular to each other and oscillate transversely to the direction of the wave's travel. The direction of the wave's travel is also the direction of energy transfer. The angle of the electric field when viewed from the direction of travel is defined as the polarization

of the wave. Polarization is a property of waves that describes the orientation by which they oscillate. It is defined as the orientation of the wave's electric field over one period. The polarization is perpendicular to the wave's direction of travel. This means that the orientation of a wave of light traveling towards an observer in the z direction can be described by two components in the x and y directions.

A wave of light that has random, sudden changes in angle and orientation with no correlation between the x and y components is defined as unpolarized. Every orientation of the wave is considered to be equally probable at any given time, so that on average, the light is symmetrical around the direction of propagation. An alternative method of imagining unpolarized light is to think of it as randomly polarized light. Partially polarized light has partial correlation between the x and y components and can be thought of as the superposition of completely polarized light with completely unpolarized light.

If the x and y components of the wave's oscillation remain constant with respect to each other and are in phase, the light is linearly polarized. When viewed from the z direction, a line is traced in the x - y plane, with the angle it makes equal to the angle of the linearly polarized light. A model of the polarization of light can be shown in Figure 18.




Figure 18: Oscillation in a wave of light viewed end-on. (A) Unpolarized light, all planes are equally probable. (B) One wave polarized light [31].

When the x and y components of a wave have the same amplitude but are 90° out of phase, the light is circularly polarized. This means that when one component of the electric field is at a maximum or minimum, the other is at zero, tracing out a circle in the x - y plane as the wave propagates. The direction of circular polarization, whether the circle is being drawn clockwise or counter-clockwise, is dependent on if the x -component is 90° ahead of the y -component or vice-versa. This is called right-hand or left-hand circular polarization. If the x and y components of the wave are out of phase and do not have the same amplitude or are not 90° out of phase, the polarization is elliptical. Examples of these three polarization states can be seen in Figure 19.




Figure 19: (left to right) Representations of linear polarization, circular polarization, and elliptical polarization. The bottom plane shows the shape traced out by the combination of the two components, the red and blue waves.

When light reflects off a surface, the plane made by the propagation direction and a vector perpendicular to the surface is known as the plane of incidence, shown in Figure 20.




Figure 20: A visualization of p-polarized and s-polarized light [32].

The component parallel to the plane is plane-polarized, or p-polarized, and the component perpendicular to the plane is sigma-polarized or s-polarized.

3.1.2 Jones Vector

The Jones vector is a mathematical method of describing the state of polarization of a beam of light. It uses the amplitude and phase of each component of the electric field to describe the state of polarization [33]. The Jones vector is only applicable to fully polarized light; it cannot characterize partial polarization.

$$\vec{E} = \begin{bmatrix} E_x(t) \\ E_y(t) \end{bmatrix} \quad (3.1)$$

$E_x(t)$ and $E_y(t)$ are the instantaneous scalar components of the electromagnetic wave \vec{E} . In complex form, the Jones vector is

$$\tilde{E} = \begin{bmatrix} E_{0x}e^{i\varphi_x} \\ E_{0y}e^{i\varphi_y} \end{bmatrix} \quad (3.2)$$

where φ_x and φ_y are the phases of each component of the wave and E_{0x} and E_{0y} are the amplitudes. For example, purely horizontally and vertically polarized light can be represented by

$$\tilde{E}_h = \begin{bmatrix} E_{0x}e^{i\varphi_x} \\ 0 \end{bmatrix} \text{ and } \tilde{E}_v = \begin{bmatrix} 0 \\ E_{0y}e^{i\varphi_y} \end{bmatrix} \quad (3.3)$$

because the y component of horizontally polarized light is 0 and the x component of vertically polarized light is 0. We know that the sum of the horizontal and vertical components of a wave of light is equal to the full intensity

$$\tilde{E} = \tilde{E}_h + \tilde{E}_v \quad (3.4)$$

If we assume, for example, that the amplitudes E_{0x} and E_{0y} are equal, and the phases φ_x and φ_y are aligned (equal), then

$$\tilde{E} = \begin{bmatrix} E_{0x}e^{i\varphi_x} \\ E_{0x}e^{i\varphi_x} \end{bmatrix} \quad (3.5)$$

If we factor out $E_{0x}e^{i\varphi_x}$, then

$$\tilde{E} = E_{0x}e^{i\varphi_x} \begin{bmatrix} 1 \\ 1 \end{bmatrix} \quad (3.6)$$

which describes linearly polarized light at 45° . If we visualize this on the elliptical plane, we would see that at the instant in time when the x and y components of the light wave reach maximum, the vector that goes through the point is at 45° above the x-axis.

To simplify the Jones vector notation, the irradiance of each component can be normalized to 1. This sacrifices the exact amplitude and phase information, but results in simpler notation. Normalization is accomplished by dividing both elements in the vector by the same scalar quantity such that the sum of the squares of the components in the vector is equal to one [33].

For example, dividing both components of Equation (3.5) by $\sqrt{2}E_{0x}e^{i\varphi_x}$ reveals the common Jones vector for 45° polarized light

$$\vec{E}_{45} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \end{bmatrix} \quad (3.7)$$

A list of common Jones vectors can be found in Appendix A.

3.1.3 Stokes Parameters

Another mathematical method of describing the polarization state of light is the Stokes vector. It has the ability to describe both fully polarized and partially polarized light. Stokes vectors can be

calculated in real time from data taken by the polarimetric camera system on a pixel by pixel basis to create a Stokes image. The Stokes images can be used to determine the state of polarization of the scene captured by the camera.

Stokes parameters are calculated from a set of four intensity measurements taken through polarizing filters. For example, the first filter could be isotropic (allow all states through), the second would horizontally polarize light, the third at 45°, and the fourth would circularly polarize light. These are the intensities I_0 , I_1 , I_2 , and I_3 . The filtered light is then measured by a polarization-insensitive device, such as a camera. The definition of the Stokes parameters is then given by

$$S_0 = 2I_0 \quad (3.8)$$

$$S_1 = 2I_1 - 2I_0 \quad (3.9)$$

$$S_2 = 2I_2 - 2I_0 \quad (3.10)$$

$$S_3 = 2I_3 - 2I_0 \quad (3.11)$$

S_0 is the overall intensity, while S_1 , S_2 , and S_3 specify the state of polarization. S_1 describes whether the polarization state trends towards horizontal (if S_1 is greater than 0) or towards vertical (if S_1 is less than 0). If S_1 is equal to 0, it may be elliptical, circular, or unpolarized [33]. S_2 describes a tendency for the polarization state to be -45°, 45°, or neither (if S_2 is greater, less than, or equal to 0). S_3 describes the tendency for the light to be circularly polarized (right handed if S_3 is greater than 0, left handed if S_3 is less than 0, or neither if S_3 is equal to 0). The combination of the Stokes parameters is called the Stokes vector, shown in Equation (3.12).

$$\begin{bmatrix} S_0 \\ S_1 \\ S_2 \\ S_3 \end{bmatrix} \quad (3.12)$$

The Stokes vector is often normalized by dividing each element by S_0 . For example, the Stokes vector for natural (unpolarized) light is

$$\begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} \quad (3.13)$$

For the purposes of the polarimetric camera, the Stokes vector can be calculated at the pixel level across the entire image. This makes it an ideal choice to be calculated from real world data sensed by the camera. Differences in the Stokes vector values between objects or areas in an image can allow a user or computer system to perform segmentation or object classification based on how they change. A list of common Stokes vectors can be found in Appendix A.

From the Stokes vector, the degree of polarization for partially polarized light can be calculated by

$$DoP = \frac{\sqrt{(S_1^2 + S_2^2 + S_3^2)}}{S_0} \quad (3.14)$$

A degree of polarization of 1 means that the light is fully polarized, while a degree of polarization of 0 means that the light is fully unpolarized. The degree of polarization is used in determining if an area of an image is polarizing light more or less than other areas of an image. Areas that are highly polarizing, such as shiny reflective surfaces, could be objects of interest for a UGV operating in a natural environment.

3.1.4 Jones and Mueller Matrices

The Jones and Muller matrices are methods of mathematically manipulating Jones and Stokes vectors by passing a beam of light through representations of optical elements, such as a polarizer.

For example, if horizontally polarized light, represented by the Jones vector

$$\tilde{E}_0 = \begin{bmatrix} 1 \\ 0 \end{bmatrix} \quad (3.15)$$

is passed through a vertical linear polarizer, represented by the Jones matrix

$$\mathcal{A} = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} \quad (3.16)$$

the final state of polarization is represented by

$$\tilde{E}_1 = \mathcal{A}\tilde{E}_0 = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \end{bmatrix} \quad (3.17)$$

with \tilde{E}_1 equaling

$$\tilde{E}_1 = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \quad (3.18)$$

This follows what we expect to occur when horizontally polarized light, which has no vertical component, is passed through a vertical linear polarizer, which only allows vertical components through, resulting in no light passing through the polarizer.

Mueller matrices manipulate the Stokes vector in a similar fashion to how the Jones matrices manipulate the Jones vector. For the same example above, horizontally polarized light represented by the Stokes vector

$$\tilde{E}_0 = \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} \quad (3.19)$$

is modified by a vertical polarizer represented by the Mueller matrix

$$\mathcal{A} = \frac{1}{2} \begin{bmatrix} 1 & -1 & 0 & 0 \\ -1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \quad (3.20)$$

with the final state of polarization, $\tilde{E}_1 = \mathcal{A}\tilde{E}_0$, equaling

$$\tilde{E}_1 = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \quad (3.21)$$

This again confirms what we know happens when horizontally polarized light is filtered with a vertical linear polarizer.

For the polarimetric camera, there are four paths of light going through the system. The four angles of polarization measured are 0° , 90° , 45° , and -45° . The -45° measurement was chosen in place of circular polarization, which does not occur in many places in nature. The general Mueller matrices modifying input light \tilde{E} for the four paths of light can be found in Table 1.

Table 1: The generic Mueller matrices for the polarimetric camera

Camera 0 (0°)	$\tilde{E}_0 = \mathcal{A}_0\tilde{E} = \frac{1}{2} \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \tilde{E}$
Camera 1 (90°)	$\tilde{E}_1 = \mathcal{A}_{90}\tilde{E} = \frac{1}{2} \begin{bmatrix} 1 & -1 & 0 & 0 \\ -1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \tilde{E}$
Camera 2 (45°)	$\tilde{E}_2 = \mathcal{A}_{45}\tilde{E} = \frac{1}{2} \begin{bmatrix} 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \tilde{E}$
Camera 3 (-45°)	$\tilde{E}_0 = \mathcal{A}_{-45}\tilde{E} = \frac{1}{2} \begin{bmatrix} 1 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 \\ -1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \tilde{E}$

A list of common Jones and Mueller matrices for various optical elements can be found in Appendix B.

3.1.5 Brewster's Angle

Sir David Brewster discovered Brewster's angle in 1815. It is a special case when dealing with light interacting with a dielectric surface at a certain angle. If the incoming light is p-polarized, it is transmitted through the surface perfectly with no reflections. If the incoming light is unpolarized, the reflected light is perfectly polarized and the transmitted light is partially polarized. A simple representation of light reflecting at Brewster's Angle is shown in Figure 21.

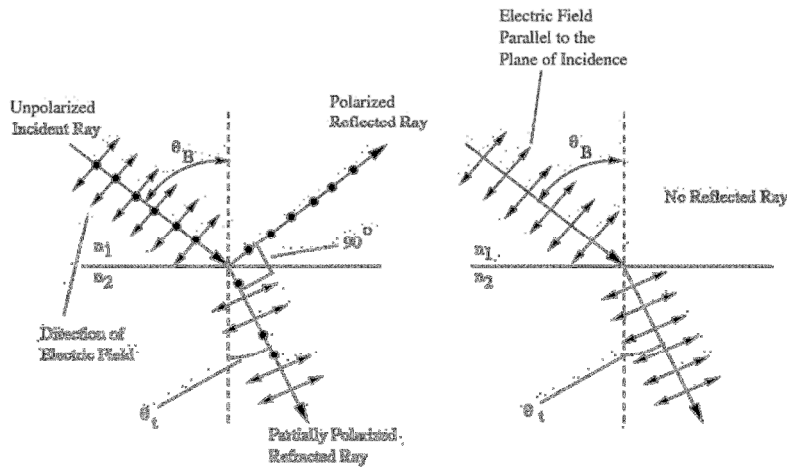


Figure 21: An illustration of the polarization of light when reflected off a dielectric surface at Brewster's Angle [34].

When light interacts with a boundary between two different indexes of refraction, some of the light is refracted and reflected. Brewster's angle is found by Brewster's Law, a special case of the Fresnel Equations:

$$\theta_B = \arctan\left(\frac{n_2}{n_1}\right) \quad (3.22)$$

where n_1 is the refractive index of the first medium the light is traveling through and n_2 is the index of the second medium. Brewster's Law states that the angle of incidence for maximum polarization depends only on the index of refraction and that maximum polarization occurs when the angle between the reflected and refracted light is 90° . Since the refractive index of a medium is a function of wavelength, Brewster's angle will also vary with wavelength. Visible light traveling through air and reflecting off glass has a Brewster's angle of approximately 56° , while visible light traveling through air and reflecting off of water is approximately 53° . For the polarimetric camera, light reflecting closer to Brewster's angle off of water or mud will be more polarized than the rest of the scene.

When the incident light is temporarily absorbed by the atoms in the dielectric material, the electrons in the atom begin to oscillate in the direction of the electric field of the refracted light. The oscillation of the electrons reemits the light and is the source of the refracted and reflected light off and through the dielectric material. The direction of the electric field of the reemitted light is the same as the direction the electrons' oscillation. At Brewster's angle, the orientation of the incoming electric field is perpendicular to the plane of incidence and parallel to the reflecting surface. Electrons oscillating in the plane of incidence are parallel to the electric field and cannot radiate in the same direction. Only electrons perpendicular to the plane of incidence can. Thus, only light oscillating perpendicular to the plane of incidence, or s-polarized light, is reflected [14], [35]. This change in the polarization of the reflection can be seen in the graph in Figure 22.




Figure 22: The reflectivity of smooth water at 20° C [36].

At 56°, the reflectivity of the s-polarized light goes to zero, leaving only the s-polarized component in the reflection. This s-polarized component is the horizontally polarized light that the polarimetric camera looks for when searching for water or mud in a scene.

3.1.6 Polarizing Filters

The most common type of polarizer in use is a Polaroid filter. Polaroid is the trademark name for the material patented by Edwin Land in 1932, originally developed to reduce glare [37]. It is a synthetic plastic sheet that contains dichroic crystals. When the sheet is stretched during manufacturing, the crystals in the plastic align along the direction it is stretched. Dichroic crystals absorb light polarized in parallel with the direction of the crystal alignment and transmit light polarized perpendicular to it. Figure 23 shows a representation of light traveling through a dichroic crystal.




Figure 23: (top) An electromagnetic wave polarized parallel to a dichroic crystal is absorbed and its intensity reduced. (bottom) An electromagnetic wave polarized perpendicularly to the crystal passes through with little change. [38]

If the orientation of the electric field, which is the same as the polarization of a wave of light, is parallel to the line of crystals, a current will begin to flow along that line, transferring energy away from the wave and reducing the intensity of the polarized light. Light polarized perpendicular to the line of crystals will produce little or no current in the crystal, which leaves the wave unchanged in that direction and able to continue through the filter. The end result is only light polarized perpendicular to the orientation of the crystals is able to pass through the filter [14], [38].

Today, the most widely used linear polarizing film is H-sheet, also invented by Edwin Land [33]. H-sheet uses hydrocarbon molecules instead of dichroic crystals, but otherwise functions in a similar fashion. Typically, the polarizing film is mounted between two glass plates for durability and the ability to orient the direction of polarization. It can also be purchased in large sheets that can be cut down to custom size and shape. For the polarimetric camera, linear polarizing film was placed in front of each camera after being cut to the correct size and orientation. The design of the filter assembly is described in detail in Section 3.2.8.

3.1.7 Beam splitters

A beam splitter is an optical element that takes in light from one direction and both transmits and reflects the beam into two separate beams, usually at right angles to each other, as shown in Figure 24. The polarimetric camera will use beam splitters to split the light four ways. Each split will then go through a linear polarizer and into a camera to be recorded.




Figure 24: An artist representation of a cube beam splitter. The image is transmitted from the left side, through the beam splitter and reflected at a 90° angle, producing two images [39].

Beam splitters can take the form of either a plate or a cube, made out of different materials depending on what wavelengths of light it needs to split. Plate beam splitters tend to be inexpensive compared to cube beam splitters. Plate beam splitters are often used for teleprompters and 3D cinema cameras, which makes economies of scale work in their favor. Cube beam splitters usually use more material than a plate beam splitter of the same size. Cube beam splitters are also often specially made to withstand high power lasers or be non-polarizing. This makes cube beam splitters more expensive compared to a similarly sized plate beam splitter. The beam splitter design for the polarimetric camera is discussed in detail in Section 3.2.6.

3.2 Detailed design and assembly

This section will discuss in detail the hardware and software design of the polarimetric camera.

3.2.1 Requirements

The main motivation for designing and building a polarimetric camera in-house was the potential low cost when compared to commercial systems. The general requirements for a prototype polarimetric camera are listed below.

1. The total camera cost must be low compared to commercially available cameras.

The initial motivation for starting the polarimetric camera design was the high cost of entry for commercial polarimetric cameras. Companies were hesitant to lend out cameras to test and were quoting between \$25,000 and \$50,000 to purchase.

2. The raw images to be used in the Stokes vector calculations must be taken simultaneously.

In order to compare the images to calculate the Stokes images, the images must be taken simultaneously. If they are not, then moving objects would appear to be in different areas of the images and would not provide useful data.

3. The raw images to be used in the Stokes vectors calculations must be registered to each other.

Related to Requirement 2, the features in each image must be registered, or aligned, to each other. The methods used for registration are discussed in section 3.2.3.

4. The camera must be easily carried by one person.
5. The camera should be easily mounted on a tripod or vehicle.

Requirements 4 and 5 are specific to the work done for this thesis. Ease of use makes collecting data easier and quicker. On a future UGV, light weight and ease of mounting would be required to ensure that adding a polarimetric camera is not a burden to the vehicle or other sensors.

6. Images should be processed at a minimum of 10Hz.

This requirement is to prove that a polarimetric camera could be used on a moving UGV. The main sensor used on many UGVs at Virginia Tech, the Velodyne 32-E LIDAR, operates at 10Hz.

3.2.2 Design Summary

The prototype polarimetric camera, seen in Figure 25, is an optically registered four camera system that calculates linear Stokes parameters to determine how light in a scene is polarized.




Figure 25: A CAD screenshot of the overall system design showing the cameras, beam splitter assembly, mounting plate, and USB hub.

The beam splitter assembly, inside the green box, splits the incoming light four ways into each camera, so that the images are all identical. The cameras are all screwed down onto the aluminum base plate and aligned so the lenses are all the same distance from their respective beam splitter. The decision was made to not design an optical system that required a forward facing lens that then focused down directly onto the separate camera sensors. This reduced the complexity of the overall design. Optics design is out of the scope of most traditional engineering coursework and would have greatly increased the amount of time needed to learn and apply the techniques needed for a prototype polarimetric camera system.

3.2.3 Image Registration

For requirement 3, there were two available methods to register images from separate cameras to each other. Registering images is defined as having the same point in a scene located at the same pixel in each image. Registering the images is a required step for calculating the Stokes parameters used in a polarimetric camera.

The first method of registering images could have been implemented in software using a feature-level detection algorithm such as the Scale-Invariant Feature Transform (SIFT) [40]. SIFT operates by detecting features in an image, such as areas of high contrast like edges or corners. The SIFT algorithm is designed to identify these features because they do not change if the images are rotated or scaled. The features are stored in an index to be compared to features in another image. The locations of matching features in both images could then be used to warp or transform the image so that the pixel locations in the images match. Figure 26 shows a pair of images that have features detected and matched using SIFT.




Figure 26: (top) A pair of images showing a rooftop scene viewed at different angles. (bottom) The same two images with matching SIFT features [41].

The two images in Figure 26 were taken with only parts of the scene showing up in both images. The SIFT algorithm matches the points it finds using features it identifies as the same in both, even after the camera perspective has changed.

There are some advantages to registering the images in software. There is no requirement to line up the cameras optically, simplifying the actual mounting and placement of the cameras. As long as all the cameras could see most of the same scene, the images could be registered. However, using feature level registration is very computer intensive and would slow down the overall frame rate of the camera. It would also dramatically increase the amount of programming needed to successfully run the camera. Also, some pixels along the borders of the image are lost and the overall image size is reduced when the images are warped and transformed. Finally, feature level registration is not guaranteed to work 100% of the time if there is a lack of features in the scene, such as large blank areas like a solid colored wall.

The second method of registering images is to optically align the cameras so that the same ray of light hits the same point on each sensor simultaneously. This method is sometimes called bore

sighting. This is achieved using beam splitters or prisms to split the light into two directions, with a camera aligned to view each image at the same distance and orientation. This method requires no extra software to achieve registration aside from potential small corrections to account for manufacturing tolerances. The downside to optically aligning the cameras is a much more complex physical design and manufacturing process to ensure registration is achieved. Additional optical hardware is also needed to split the light, increasing the total cost of the camera system. Adjustments to the image alignment would be further tuned in software by translating and rotating the images if necessary to achieve a proper alignment.

In the end, the decision was made to pursue optical registration over software based registration. Optical registration, once achieved, will work regardless of what the contents of the scene hold. Software-based registration could fail if the algorithm does not detect enough feature points to make an accurate transformation. It also simplifies the program needed to capture and process the images, which keeps the frame rate of the camera higher and enables it to operate in real time.

3.2.4 Optical Components and Camera Details

The following section will describe in detail the beam splitters chosen to be the heart of the polarimetric camera, as well as the Point Grey Firefly MVs that acted as the main sensors. It will also describe the triggering process that ensured the cameras had synchronized images.

3.2.4.1 Beam splitter selection

The beam splitters used in the prototype camera had some key requirements.

1. The beam splitters must reflect and transmit the incoming light without changing its polarization state.
2. The beam splitters must be large enough to encompass the field of view of the cameras.
3. The beam splitters must reflect and transmit light equally.

For requirement 1, this required a non-polarizing beam splitter. Traditional plate beam splitters polarize the light upon reflection, which would destroy any of the original polarization data coming into the cameras. Non-polarizing beam splitters are available in cube form in various sizes. With a 16mm lens on each camera, a 20mm cube beam splitter is just large enough to encompass the field of view. This was calculated using the simplified equation for the field of view of a lens

$$\alpha = 2 \arctan \left(\frac{d}{2f} \right) \quad (3.23)$$

where d is the size of the sensor in the measured direction, f is the focal length of the lens, and α is the angle of view in the measured direction. Using CAD software, a model of the field of view was constructed and used in positioning the cameras relative to the beam splitters so that their fields of view did not exit the beam splitter assembly through the side faces.

The beam splitter chosen was a ThorLabs 50:50 non-polarizing cube beam splitter, shown in Figure 27.



Figure 27: ThorLabs BS016 50:50 Non-Polarizing Beam splitting cube. [42]

The BS016 beam splitter is made of N-BK7 optical glass, which has a transmission range from 350-2500nm. It has an anti-reflective coating that is optimized for visible ranges between 400-700nm and has a purchase price of \$181.50 when purchased in quantities of one. [42].

3.2.4.2 Camera Selection

A large portion of the total cost of the system is machine vision cameras. There were several requirements for the cameras in order to meet some of the general requirements detailed in section 3.2.1.

1. The cameras must have an external trigger.
2. The cameras must come with a C++ API.
3. The cameras must have a standard interchangeable lens mount.
4. The cameras must have a relatively low cost.

External triggering and the C++ API allow images to be captured at exactly the same moment in time. The images would then be read by a custom computer program for further processing. An interchangeable standard lens mount would allow for a range of focal lengths to be tested to determine what the best one for prototyping purposes would be without the need for complicated optics calculations or modeling. Finally, low camera cost keeps the overall polarimetric camera cost low.

The Point Grey Firefly MV, model FMVU-03MTM/C, seen in Figure 28, was the machine vision camera selected.



Figure 28: The Point Grey Firefly MV machine vision camera [43].


The Firefly MV boasts a 1/3" CMOS sensor capable of outputting 640x480 pixel images at up to 30 frames per second in a free running mode, or 15 frames per second in an externally triggered mode. The CMOS sensor has a global shutter, meaning the entire sensor is read at once, instead of line by line, avoiding wobble or warping if the camera encounters rapid movement or vibration. The Firefly can accept any C or CS mount lens.

The Firefly MV both receives power and transmits data over a single USB 2.0 cable. The cameras are all plugged into a powered USB hub that then has a single USB cable going to the computer. This simplifies connecting to the polarimetric camera and reduces the amount of cables needed. Finally, when purchased in quantities of 5, the Firefly MVs cost \$200 each. In quantities of one, price goes up to \$275 each.

3.2.4.3 Triggering

Requirement 2 from section 3.2.1 states that all the cameras in the system must capture images simultaneously. The proper way to achieve this is to synchronize the cameras using a hardware trigger.

A hardware trigger is an electrical signal that is sent to a digital input line on the camera to start a new frame. The signal is usually a step change from 0 to 3.3V that is pulsed at the desired frame rate. A diagram of the external trigger timing from the Firefly MV manual is shown in Figure 29.



	FFMV-03M2/FMVU-03MT	FMVU-13S2
1	Min 1 μ s	Min 1 μ s
2	less than 10 μ s	Varies according to shutter speed; approximately the difference between frame rate and shutter speed.
3	Exposure time	Exposure time
4	1 ms	1 -2 horizontal line scans; horizontal line frequency can be read from register 0x1AF4.
5	30 ms (30 FPS)	30 ms (30 FPS)

Figure 29: Point Grey Firefly MV external trigger timing characteristics [43].

The external trigger signal at step 1 starts the camera exposure time at step 2. After the image is finished exposing during step 3, the camera begins data transfer at steps 4 and 5. Once data transfer is complete, a new trigger signal can be accepted to start the next frame.

Typically, a hardware trigger is sent from a printed circuit board (PCB). For the prototype polarimetric camera, the level of simplicity needed for triggering meant that an off the shelf solution could be used in place of a custom designed board. The Texas Instruments MSP430 LaunchPad, seen in Figure 30, was chosen to supply the hardware trigger signal.



Figure 30: A top down view of the MSP430 LaunchPad development board [44].

The MSP430 LaunchPad is a developer board made by Texas Instruments as a method of introducing their line of MSP430 processors to a wide range of potential developers by offering them at low cost. It offers several digital inputs and outputs that can be configured to be GPIO, or communication lines such as RS232. It also has a reset and interrupt button, as well as two LED indicators. The LaunchPad is programmed via a USB interface, which also delivers power.

The USB power interface is ideal for the prototype polarimetric camera as it will already have a powered USB hub inside delivering power to the Firefly MV cameras.

The LaunchPad is programmed in C using TI's Code Composer Studio development environment. A very simple program was written to trigger the cameras at a rate of 15Hz, which is the maximum rate that a Firefly MV can be externally triggered. The source code for the program can be found in Appendix C.

3.2.5 Base Platform

The main mounting platform for the cameras and beam splitters is a plate that is CNC machined out of 6061 aluminum, shown in Figure 31. The plate also has a raised platform for the beam splitter assembly.




Figure 31: A CAD screenshot of the aluminum base plate.

The beam splitter platform is 7mm tall, placing the vertical center of the beam splitters at the same height as the center of the Firefly MV camera sensors.

The front of the camera is facing towards the upper right corner in Figure 31. The aluminum plate has M2 clearance holes drilled for 3 Firefly MV cameras, spaced around three edges of the beam splitter platform. Around the edges are 1/8in. clearance holes to secure the plate into the bottom of a plastic enclosure. The last sets of holes drilled are tapped 1/8in. holes for securing the beam splitter clamp and box to the plate. The majority of the plate area that holds the cameras and USB hub mounted is 0.25in. thick.

The base plate was milled out of 6061 aluminum by Metal Processing, Inc. in Radford, VA. The cost associated with a third party doing the machining was determined to be less than milling it in-house.

3.2.6 Beam Splitter Layout

The beam splitters selected in section 3.1.7 were ThorLabs 20mm beam splitters. A diagram of how the beam splitters are laid out is shown in Figure 32 with the front of the camera is facing to the right.

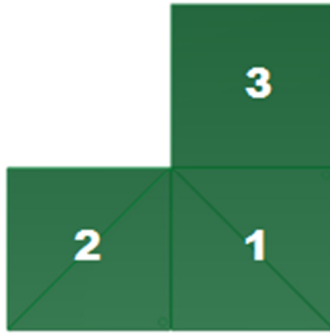


Figure 32: A top-down view of the beam splitters, labeled 1 through 3.

In this orientation, the incoming light enters from the right into beam splitter #1. With the beam splitters in this configuration, there is not enough room to have all four cameras in the same plane as the beam splitters without at least one camera partially blocking the others or the path of the incoming light. Thus, the beam splitters are oriented so that the light travels in the path shown in Figure 33.




Figure 33: A top-down view of the path light takes inside the beam splitter assembly. The dark blue arrow is the entrance face of the assembly. The red arrows show how the light is split inside. The light blue lines and square are where the Firefly MVs are located.

The incoming light, shown in dark blue, enters from the right into beam splitter #1, which reflects half the light into beam splitter #3 and transmits the other half into beam splitter #2. In beam splitter #2, the light is split again and enters two Firefly MVs, which are represented by light blue planes. For beam splitter #3, space constraints required a different orientation. Beam splitter #3 is set so the reflected portion is directed upwards instead of to a side, shown in Figure 34.




Figure 34: A front view of the beam splitter assembly (Not visible: beam splitter #2).

The transmitted portion continues through beam splitter #3 and into the third Firefly MV. This allows for the fourth Firefly MV camera to be placed above the beam splitter assembly without interfering with the other three Firefly MVs or the incoming light.

3.2.7 Beam Splitter Mount Design and Results

The beam splitters from ThorLabs were unmounted. Thus, a method for securing them to the baseplate had to be devised. Two different methods are explored in this section, a clamp method and an adhesive method. In the end, the adhesive method was used with good results.




Figure 35: A CAD screenshot of the beam splitter clamp method (highlighted in blue).

The clamp method, shown in Figure 35, would press the beam splitter assembly into the raised lips on the sides of the beam splitter platform. The clamp piece would be pushed into the beam splitters and then bolted down, keeping the clamping force on the sides of the cubes.

The alternative to the clamp method is to attach the beam splitter directly to the aluminum using an adhesive. In the auto industry, glass and metal are most commonly adhered in the rear view mirror assembly. The adhesive must withstand daily use by the driver, UV radiation, large temperature ranges and swings, as well as vibration and shock. These adhesives are inexpensive and readily available in any hardware or auto parts store.

An adhesive method would simplify the overall design, since it does not require any additional components or assembly to work. Commercial adhesives also have a very high holding strength, which greatly surpasses the actual force needed to hold the beam splitters in. Commercially available mounted cube beam splitters are held in their mounts with adhesive, not mechanically. However, an adhesive method would be very difficult, if not impossible, to undo or correct once set. This would make any mistakes in assembly permanent.

In practice, the beam splitter clamp system did not work as designed. Due to small manufacturing tolerances in the beam splitters, force could only be applied normal to one face of

the beam splitter assembly. This secured beam splitters #2 and #3, but could not secure beam splitter #1. Thus, the epoxy method had to be used. Devcon 2 Ton Epoxy was the adhesive used to secure the beam splitters to the aluminum. Tests on scrap glass and aluminum showed it had excellent holding strength and dried clear.

The beam splitters were glued in one at a time to ensure they were properly seated against the raised lip of the beam splitter platform and square to each other. An image of the beam splitter assembly after it was placed onto the base plate is shown in Figure 36.




Figure 36: An image of the beam splitters during assembly of the camera.

Index matching fluid from Cargille Labs was applied to the surfaces between each beam splitter before the faces were matched up. The index matching fluid reduces the number of interfaces where the index of refraction would change. This reduces the possibility of internal reflections occurring inside the beam splitter assembly.

3.2.8 Polarizing Filters and Holder

The polarimetric camera requires that each Firefly MV has its own dedicated polarizing filter to look through in order to calculate the differences in polarization needed. This requires a unique filter mounting solution that fits into the overall design without obstructing any beam splitter paths or obstructing camera fields of view. A CAD image of the filter holder can be seen in Figure 37.




Figure 37: A CAD screenshot of the filter holder box. The first opening is on the bottom right of the box.

The filter holder serves two purposes. First, it holds the polarization filters in their correct orientation in front of each Firefly MV camera. This is accomplished by including a small 2mm inset on the inner side of each opening. From the outside, the opening has a 2mm smaller diameter than the filters.. This gives the filter an inset area to fit snugly into. For the linear polarizers, alignment is carefully performed by hand to insert the filter at the correct angle for each camera. The filters would be secured into place with a small amount of adhesive to keep them from moving.

The second function of the filter holder is to stop any stray light from entering the beam splitter assembly. For the best image quality, light can only enter from the front of the beam splitter assembly, through the opening on the bottom right on Figure 37. With the camera lenses very close to the other openings in the filter holder, little stray light can enter the assembly. Painting it black further reduces any chances of light reflecting inside the filter holder and into the beam splitter assembly.

To manufacture the filter holder, an additive manufacturing method was used. Because complexity is “free” with additive manufacturing, the complex shape of the filter holder was rendered moot. The filter holder was printed on a Stratasys 768 fused deposition machine courtesy of the DREAMS Lab at Virginia Tech. If the filter holder was printed commercially, it would cost about \$31. An image of the filter holder after it was printed is shown in Figure 38.




Figure 38: An image of the filter holder after it was printed.

The filter holder was spray painted a flat black to reduce the light levels inside further. After the filter holder was printed, there were some issues with the design that were observed. While additive manufacturing made it easy to print a filter holder box that would otherwise be difficult or expensive to machine, the tolerances were difficult to get correct. Tolerances varied depending on the print direction and made it difficult to size the diameters of the insets where the unmounted glass polarizers would sit. However, a simple substitute was implemented by swapping the unmounted polarizing filters with linear polarizing film. The film was simply adhered to the face of the box over each opening at the correct orientation for each camera to look through.

To correctly determine the state of polarization using the Stokes parameters, the polarizing filters in front of each camera had to be correctly oriented. This involved both characterizing the filters as well as placing them in the correct orientation in front of each camera. The list of cameras and the filtered light each one measures can be found in Table 2.

Table 2: The polarizer-camera pairs in the polarimetric camera

Firefly MV camera number	Intensity Measured	Polarization State Measured
Camera 0	I_0	0° Polarization (Horizontal)
Camera 1	I_1	90° Polarization (Vertical)
Camera 2	I_2	45° Polarization
Camera 3	I_3	-45° Polarization

The intensity measured by each camera is the same intensity that is used to calculate the Stokes vector in Equation (3.12). The Stokes vector calculated from these measurements for this thesis is

$$\begin{bmatrix} S_0 \\ S_1 \\ S_2 \\ S_3 \end{bmatrix} = \begin{bmatrix} I_0 + I_1 \\ I_0 - I_1 \\ 2I_2 - (I_0 + I_1) \\ 2I_3 - (I_0 + I_1) \end{bmatrix} \quad (3.24)$$

S_0 is the overall intensity of the scene, S_1 is the difference between horizontal and vertical polarization, S_2 is the difference between 45° polarization and the overall intensity, and S_3 is the difference between -45° polarization and the overall intensity.

The decision to not measure circularly polarized light was made because there are very few instances in nature where naturally occurring circular polarization is found. For the purposes of classifying objects for a UGV, linear polarization is the primary modality of interest. This means for the purposes of this camera, the degree of polarization in Equation (3.14) would actually measure the degree of linear polarization, since there is no circular polarization data available.

To verify that the 45° and -45° polarizers on Camera 2 and Camera 3 were working as intended, the system was pointed at a -45° polarized light source, an Apple MacBook Pro laptop screen. The 45° polarization of the screen was discovered during the examination of a linear polarizer with the direction of filtering marked from the factory. When the filter was placed on the screen and turned at a 45° angle, it went black.




Figure 39: A diagram showing how light from an LCD screen with -45° polarization is filtered by linear polarizers placed at different angles.

After the filter was rotated 90° , the filtered and unfiltered portions of the monitor had similar intensities, meaning that the polarization angle of the filter had matched the polarization angle of the laptop screen at -45° . When the filter was placed at a 0° and 90° orientation, the screen intensity was about half.

The camera was first pointed at an unpolarized light source, a blank cinderblock wall. The exposures were matched so that the average intensities of each camera calculated in MATLAB differed by less than 0.02 on a scale of 0 to 1. A histogram of all four cameras can be seen in Figure 40.

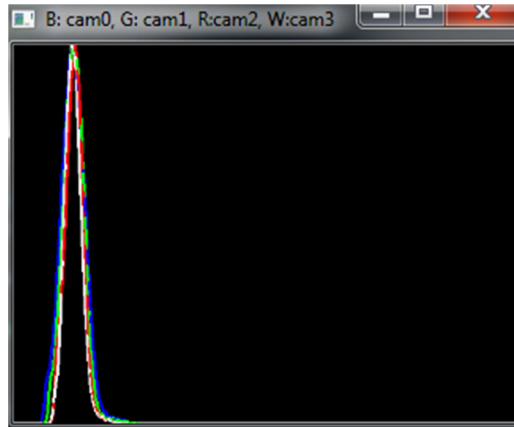


Figure 40: A histogram of all four cameras showing equal intensities when pointed at a blank wall.

With the same settings, the camera was pointed at the laptop monitor in a horizontal orientation. Images from all four cameras can be seen in Figure 41.



Figure 41: Raw Firefly MV images of a Macbook Pro laptop screen displaying an image. (top left) Camera 0, (top right) Camera 1, (bottom left) Camera 2, (bottom right) Camera 3.

As expected, the -45° linear polarizer on Camera 3, which is at the same polarization angle of the laptop screen, allows the light to pass through unfiltered. The 45° polarizer on Camera 2, which is oriented 90° offset to the polarization of the laptop screen, filters the majority of the light out. The mean intensity of the image in Camera 2 is 0.0257, while the mean intensity of the image in Camera 3 is 0.1165. With theoretical ideal polarizers and beam splitter optics, the crossed polarizer in Camera 2 should result in zero intensity.

The horizontal and vertical images from Camera 0 and Camera 1 in Figure 41 have mean intensities of 0.0515 and 0.0608, respectively. This is what is expected when looking at a -45° polarized light source. The intensities of the horizontal and vertical components are roughly equal and are about half the intensity of the -45° polarized filter. We can verify this by placing the camera at a 45° angle relative to the screen, shown in Figure 42.




Figure 42: Raw images of a Macbook Pro laptop screen taken at a 45° angle. (top left) Camera 0, (top right) Camera 1, (bottom left) Camera 2, (bottom right) Camera 3.

At a 45° angle, we can see that the image intensity has dropped in Camera 1, which is the vertically polarized sensor. This matches what we expect to happen: the vertically polarized filter is at a 90° offset to the -45° polarization of the monitor, allowing very little light through. The mean intensity of Camera 2 and Camera 3 are roughly equal, at 0.0626 and 0.0720 respectively.

3.2.9 Software

The software for the prototype polarization camera has several main requirements.

- The software must acquire images from all four cameras simultaneously.
- The software will compute the linear stokes parameters.
- The software will correct for small image alignment errors.
- The software will display raw and processed images in real time.
- The software will save raw and processed images to disk.

Point Grey includes the FlyCapture2 C++ SDK for use with the Firefly MV cameras. The FlyCapture2 SDK includes functions that can acquire data and control camera parameters such as triggering modes or shutter speed.

The OpenCV C++ library was selected for image processing. OpenCV is the industry standard image library used for both commercial and research applications. OpenCV is an open source C++ library to develop software for PC, Mac, and Linux, as well as the Android and iOS mobile platforms. It has over 2500 built-in functions for image processing and machine learning.

The QT Creator IDE was chosen to integrate the FlyCapture2 SDK and the OpenCV libraries. QT Creator also has an integrated GUI development kit that can be set up to display OpenCV image types. QT Creator greatly reduced the time needed to set up a usable GUI as well as integrate the FlyCapture2 and OpenCV libraries. The GUI developed for the polarimetric camera can be seen in Figure 43.

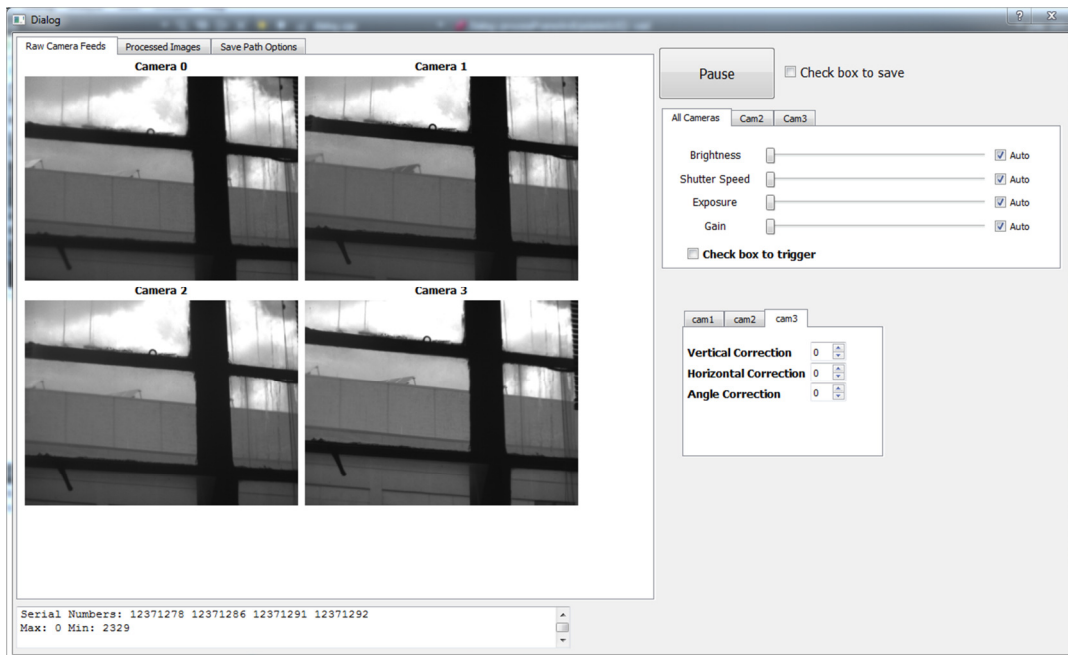


Figure 43: A screenshot of the software GUI for the polarimetric camera.

The software developed for the camera has several main functions. It reads in images from the cameras and displays the raw data to the user. The user also has the ability to manually control camera parameters such as shutter speed and gain. By default, the cameras control these parameters automatically when first powered up. To aid in adjusting camera exposure parameters, the software calculates the mean intensity of each camera in real time and displays them to the user. The software also calculates the Stokes parameters and the degree of polarization. It can show the final processed data to the user. Finally, both raw images and processed images can be saved to disk in real time as the cameras capture new frames.

Tolerances required to manufacture and assemble the baseplate meant that a perfectly aligned image with no corrections needed is unlikely. Therefore, the final image alignment was completed by the camera software. The software image alignment correction only needs to be

done once, as the cameras and beam splitter would not move relative to each other once they are both secured.

The source code for the polarimetric camera software developed for this thesis can be found in Appendix D.

4 Experiments

4.0 Summary of Chapter

The first section of this chapter will discuss the verification of data produced by the polarimetric camera. This includes image alignment between cameras, which is required to produce the Stokes images. It will then discuss an observed intensity difference between the Firefly cameras as they look into different outputs of the beam splitter assembly. Finally, the verification section will review data and calculated Stokes images of a controlled polarized light source.

The second section of chapter 4 will take the polarimetric camera outdoors to observe natural and man-made environments. The natural scenes are focused on standing water and muddy paths after a heavy rain storm, as well as running water on the Virginia Tech campus. Man-made scenes observed include roads, buildings, and vehicles.

The step-by-step process used to collect data with the polarimetric camera can be found in Appendix E.

4.1 Indoor verification and tests

4.1.1 Image alignment between Firefly cameras

Accurate alignment is a requirement for the polarimetric camera system to function properly. There were several sources of error that contributed to the amount of correction needed in software.

The most common source of error in aligning the Firefly MV cameras was in the clearance holes needed to secure the cameras to the base plate with screws. Since a certain amount of tolerance is needed to allow the screws to slide smoothly through the plate and into the cameras, exact placement of the cameras relative to each other is not possible. In addition to the mounting hole error, the aluminum angle used as the bracket was found to be slightly acute and not at a true 90° angle. This caused a larger vertical error in Camera 3 compared to Camera 1 and Camera 2. The final source of error was discovered in the lenses used on the Firefly MVs. The inexpensive CS-mount lenses purchased for the cameras shifted the images slightly as the focus ring was rotated. However, the lenses tended to shift the image equally when all the lenses were set to the same focus point. In the future, this could be avoided with higher quality lenses.

Initial alignment of the cameras was done by hand, one camera at a time. The reference camera was assumed to be Camera 0, with the other three cameras attempting to align to it. The image from Camera 1 was subtracted from Camera 0. The resulting image from the subtraction would clearly show any horizontal or vertical errors when looking at objects in the scene with horizontal or vertical features. Figure 44 shows a sample scene taken with Camera 0 and Camera 1. The window frame, which contains the horizontal and vertical features that will be examined, is approximately 13 feet away from the camera.




Figure 44: A typical scene looking from Randolph Hall towards Whitemore Hall at Virginia Tech.

In this image, the dark window frames in the foreground clearly stand out from the brighter building in the background, making it ideal for alignment. The subtracted image is shown in Figure 45.




Figure 45: The subtracted image from Camera 0 and Camera 1.

The white areas in Figure 45 show where the images from Camera 0 and Camera 1 are out of alignment. In order to align the images, they had to be shifted in the X and Y directions. This was accomplished by placing the original 640x480 image array centered in a 744x560 blank image array. This places a buffer area around the images that lets the user move an image up to 40 pixels in any direction. An example image with the blank border is shown in Figure 46.



Figure 46: A raw image placed into a larger image array.

The border gives a buffer area for the image to move into as it is shifted around during alignment. Each camera is shifted relative to the image from Camera 0. After the image was

shifted to align the edges, the image in Figure 45 went dark, with no white areas showing an overlap.

To determine the amount of pixel shift needed for objects at varying distances, the camera was pointed at a standard grid target commonly used in characterizing stereo camera rigs or lenses. The sharp, contrasting horizontal and vertical lines produced by the checkerboard provides a good target to identify areas of misalignment. An image of the test setup is shown in Figure 47.



Figure 47: An image showing the test setup used to determine pixel shift at various distances.

A sample image of the grid as observed by Camera 0 and Cameras 1 at five feet is shown in Figure 48.




Figure 48: The calibration grid as observed by Camera 0 (left) and Camera 1 (right) at a distance of five feet.

Again, the images were subtracted and displayed in real time as pixel shifts and affine transforms were introduced to line up the horizontal and vertical edges of the grid. The subtracted image before alignment is shown in Figure 49.




Figure 49: The subtracted image from Camera 0 and Camera 1 of the grid shown in Figure 48.

The white gaps are areas of misalignment that are present before the software alignment. After alignment, the white gaps disappear, leaving behind a dark image. The calibration grid target was placed at distances of 5, 10, 15, 20, 25, 30, 40, and 50 feet away from the camera. The program was reset at each distance to start the calibration process with zero pixel shifts in any direction. A graph showing each camera’s pixel shift in the X and Y directions is shown in Figure 50.

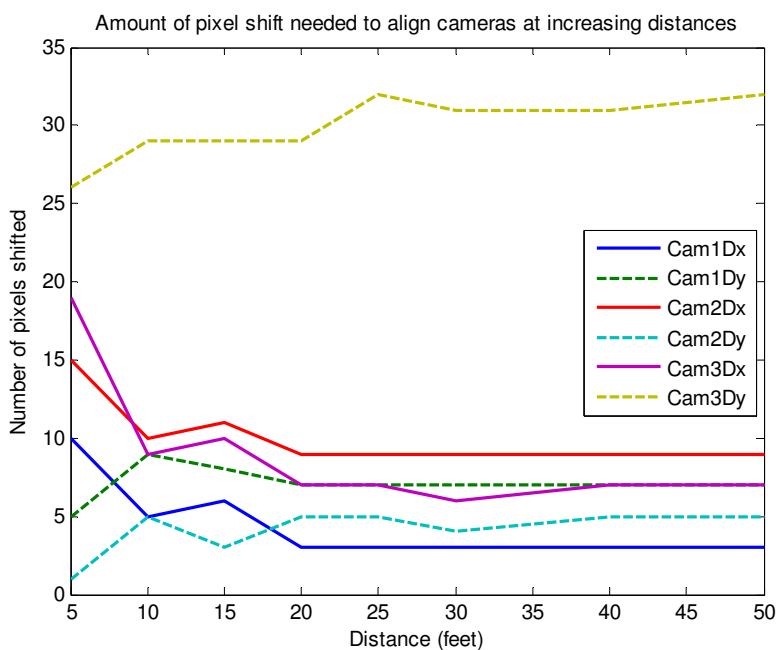


Figure 50: A graph showing the relationship between the amount of pixels Cameras 1-3 had to be shifted to align with Camera 0.

Figure 50 shows the relationship between object distance and the amount of shifting each image has to go through to be aligned with one another. The most important takeaway from Figure 50 is that all the cameras reached a “steady state” where the amount of pixel shift stayed the same or within one pixel. At about the same distance, the focus distance on the lenses reached infinity, meaning that all objects beyond about 25 feet came into focus. Any objects beyond this point will be in focus.

The amount of physical space that is observed by a single pixel at 50 feet is much larger than at 5 feet, so less correction was needed. Eventually, the target grid would converge into a single pixel

located at the same point in all four cameras. This means that in a future version of the polarimetric camera, careful lens and camera selection can be combined with an assumption that any targets of interest would be at some minimum distance away from the sensor. This would mean only one initial calibration would be needed which would hold true for objects at that minimum distance and beyond. This calibration process only takes a few minutes and can be done with little training beforehand.

A similar design decision that is commonly made for cameras on a UGV is to set the lens focal point to infinity focus. This would ensure that anything the camera looks at beyond its infinite focus location would be in focus, which is useful for a UGV that is looking several tens of meters ahead for objects.

4.1.2 Intensity differences between cameras

When testing first started, an apparent difference in intensity was immediately obvious between the raw images when the cameras were set to identical exposures. To quantify the difference, the camera was pointed at an unpolarized light source, a standard CFL light bulb. The raw images are shown in Figure 51.




Figure 51: A set of raw images of a CFL light bulb taken with the prototype polarimetric camera.

The Firefly MV cameras were commanded to their fastest shutter speed, lowest gain and exposure settings, and maximum brightness setting. Camera 3 immediately stands out as much darker than the other three cameras. Camera 2 is slightly darker than Camera 0 and Camera 1, which are indistinguishable between each other.

Using MATLAB, the mean intensity value of each image was calculated, listed in Table 3. The values range from 0 to 1 due to how MATLAB handles the double data type.

Table 3: Mean intensities of each camera of the scene in Figure 51.

Camera	Mean Intensity	Percent Difference From Max
Camera 0	0.0559	5.6
Camera 1	0.0592	0
Camera 2	0.0458	22.7
Camera 3	0.0278	53.1

The mean values verify the visual observations of each camera, with a small difference between Camera 0 and Camera 1, and much larger differences between Camera 2 and Camera 3. A second image set is shown in Figure 52.




Figure 52: A set of images taken looking from inside Randolph Hall towards Whittemore Hall at Virginia Tech.

The images in Figure 52 were taken with brightness, exposure, and gain settings set to their minimum values and shutter speed set to 0.625ms. The differences in average intensity are listed in Table 4.

Table 4: The mean intensities of the images in Figure 52.

Camera	Mean Intensity	Percent Difference From Max
Camera 0	0.3848	0
Camera 1	0.3639	5.43
Camera 2	0.2953	23.26
Camera 3	0.0396	89.70

The difference between Camera 0, Camera 1, and Camera 2 were all similar to the values in Table 3, with the main difference being the reversal of Camera 0 and Camera 1 as the brightest in the group. Camera 3 was again the darkest image, this time 89% darker than the brightest image.

The intensity difference is most likely due to the beam splitter assembly used in the camera. The two darker cameras, Camera 2 and Camera 3, both observe light that has been first reflected in beam splitter #1 before splitting in beam splitter #3.

As shown in Figure 53, the light reflected from beam splitter #1 enters beam splitter #3, where it is transmitted into Camera 2 and reflected into Camera 3, located above it.




Figure 53: A top-down view of the beam splitters, labeled 1 through 3.

Camera 0 and Camera 1 observe light that is first transmitted through beam splitter #1 before being split in beam splitter #2. In addition, Camera 3 looks downward into the aluminum base plate, while the other three do not.

To ensure that the cameras themselves were not faulty, Camera 2 and Camera 3 had their locations swapped. No intensity difference was observed and the cameras were switched back. The intensity drop was also observed without the filter holder box in place, indicating that the box and polarizing filters were also not the cause of the intensity difference.

The beam splitters are non-polarizing 50:50 beam splitting cubes from ThorLabs. They have an anti-reflective coating on them that is effective from 400-700nm. The transmission graph for the beam splitter is shown in Figure 54.

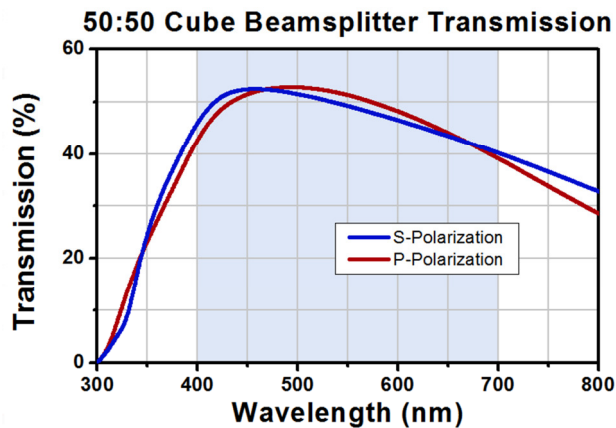


Figure 54: The transmission graph for the ThorLabs 50:50 non-polarizing cube beam splitter [42].

Looking at the transmission graph, the percent transmitted drops off as wavelength increases. This means that as wavelength increases, more of the light should be reflected versus transmitted. However, in tests, the drop off does not account for the dramatic drop off in intensity in the images from the Firefly cameras.

To characterize how the beam splitters change the wavelengths of light that pass through them, an Ocean Optics visible spectrometer was aimed into the beam splitter assembly where Camera 3 and Camera 0 are located. A white LED light source was aimed into the front of the camera, which has a spectral response shown in Figure 55.

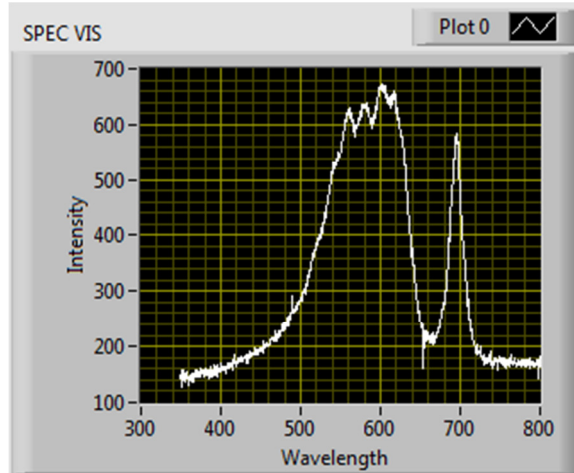


Figure 55: The spectral response of the white LED aimed into the beam splitter assembly.

The spectral response observed through the beam splitter assembly is shown in Figure 56.

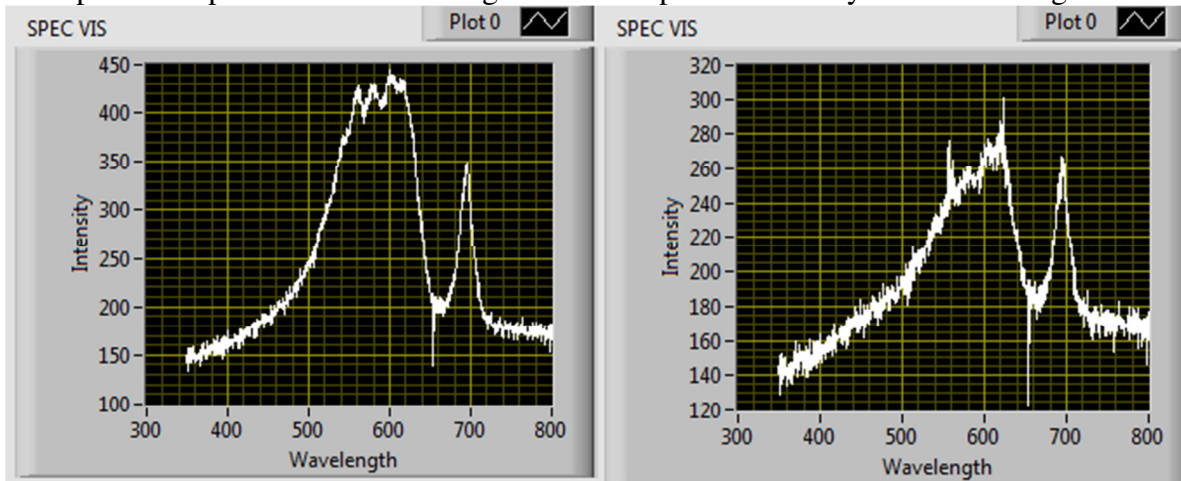


Figure 56: The spectral response of a white LED observed from the point of view of Camera 0 (left) and Camera 3 (right).

While there is a difference in the intensity of the spectrometer reading between Camera 0 and Camera 3, there is no strong difference in the wavelengths that each camera is observing.

For the purposes of observing the polarization states of various objects or scenes, the camera could still be used as long as the intensities of each Firefly camera were matched beforehand. This was accomplished by observing a non-polarizing scene, such as a blank wall indoors or by

pointing the camera straight down at the ground, and adjusting exposure levels of each camera so that the mean intensities of each camera's image were equal. The camera was then able to be pointed at the scene of interest to measure its polarization characteristics.

4.1.3 Stokes images of a polarized light source

Using a similar technique applied in Section 3.2.8, we can verify that the Stokes parameters are calculated as we expected. The same set of images of an LCD screen with -45° polarization is shown in Figure 57. The images were taken after camera intensities were equalized on a non-polarized source and aligned. The borders of some images have areas where the alignment process overlapped an edge with a blank border, resulting in incorrect data that would normally be ignored.



Figure 57: A set of raw images of a -45° polarized LCD monitor. (top left) Camera 0, (top right) Camera 1, (bottom left) Camera 2, (bottom right) Camera 3.

We can see again that the image from Camera 2, which has a 45° polarizer over it, is very dark compared to the other three cameras. This is because the -45° polarizer over Camera 2 is at a 90° angle to the 45° polarization of the monitor, effectively filtering out all the light from the monitor. Camera 3 has the highest intensity because its -45° polarizing filter matches the polarization angle of the monitor at -45° . The calculated Stokes images are shown in Figure 58.



Figure 58: Stokes images of a -45° polarized LCD monitor. The images are shown on a scale of 0 to 1, truncating any negative values for visualization purposes. (top left) S_0 , (top right) S_1 , (bottom left) S_2 , (bottom right) S_3 .

Looking at the Stokes images, there are several conclusions that can be made. First, we can see that overall value of S_1 is very close to zero. This is expected because the intensities of Camera 0 and Camera 1, the horizontal and vertical polarized cameras, are about equal because they are each capturing about half of the intensity of the -45° polarized light. Since their intensities are about equal, their subtracted values come close to cancelling out. Using MATLAB, the calculated mean intensity of S_1 is -0.0093 .

Second, S_2 also appears dark. S_2 is calculated by $2I_2 - S_0$. Because the mean value of Camera 2 is very close to 0, S_2 drops into the negatives when S_0 is subtracted, with a mean value of -0.0608 . This correlates with the fact that the camera was looking at a -45° strongly polarized source of light. A larger negative value in S_2 means that the light entering the camera has very little polarization at a 45° angle.

Third, S_3 has a very high overall value, calculated by $2I_3 - S_0$. Since Camera 3 had the highest intensity out of all the cameras, S_3 returns a positive overall value, with a mean of 0.1207 . This is expected because the polarization of the monitor is at a -45° , which is the same angle that Camera 3's filter is set at. The positive value in S_3 correlates with a strong -45° polarization, much like how a negative value in S_2 means the polarization is not at a 45° angle.

The degree of polarization, listed in Equation (3.14), can also be calculated from the Stokes parameters. To show this, the camera was aimed to include both polarized and non-polarized sources, an LCD monitor and a cinderblock wall. The raw images are shown below. The monitor was displaying a blank white area to make visualization easier.




Figure 59: A set of raw images of a -45° polarized LCD monitor on the left and a cinderblock wall on the right. (top left) Camera 0, (top right) Camera 1, (bottom left) Camera 2, (bottom right) Camera 3.

As before, the -45° polarized image from Camera 3 is the brightest and the 45° polarized image from Camera 2 is the darkest, with the intensities of the horizontal and vertical images from Camera 0 and Camera 1 in between. The Stokes images calculated from these images are shown in Figure 60.



Figure 60: Stokes images of a -45° polarized LCD monitor and cinderblock wall. (top left) S0, (top right) S1, (bottom left) S2, (bottom right) S3.

Looking at the Stokes images, S0 and S3 have a higher intensity than S1 and S2, which are close to zero intensity. The degree of polarization image is shown in Figure 61.




Figure 61: The degree of polarization image of an LCD monitor on the left and a cinderblock wall on the right.

As we expect, the area of the image that contains the LCD monitor has a high degree of polarization compared to the cinderblock wall and the frame of the monitor. The average degree of polarization value of the LCD screen is 0.9225, compared with the wall at 0.1185. This confirms that the light from the LCD monitor is highly polarized light and the light coming from the wall is not.

4.2 Polarization in the outdoors

This section will review scenes and images taken outdoors at various time of day. The polarization state of many objects depends on the angle of the camera relative to the position of the sun in the sky. Regardless, polarization remains an intriguing source of information that could potentially be applied to a future unmanned ground vehicle in a system that uses methods to segment images from each other using intensities or more complex techniques such as machine learning algorithms that can pick up more subtle differences between images and use prior knowledge to make its decisions.

4.2.1 Mud and Still Water

To observe standing water in a natural environment, data was taken at the Drillfield on Virginia Tech's campus after 36 hours of heavy rains had swept through the area. Several places on the Drillfield had standing water more than an inch deep, while other places were muddy from students walking to class through the field. The sun was at an elevation of about 37° above the horizon in the western part of the sky and was not covered by clouds.

A muddy path surrounded by grass was observed from various directions relative to the sun. The path and the surrounding grass are shown in Figure 62 below.




Figure 62: An image of the type of muddy path on the Drillfield at Virginia Tech after heavy rain.

The degree of polarization of the muddy path taken from the same point of view as Figure 62 is shown in Figure 63.




Figure 63: The S0 Stokes image (left), degree of polarization (center), and polarization contrast (right) of the muddy path.

In this figure, the sun is to the left of the image. We can clearly see that the muddy area is brighter than much of the surrounding grass. There is also a high polarized response in the grass towards the bottom of the image, where it was still wet from the rain. The polarization contrast image of the same scene is darker than the degree of polarization, but the muddy path is still contrasted with the grass.

After the rain storm, there were also several areas of standing water. An example image is shown in Figure 64.




Figure 64: S0 image (left), degree of polarization (center), and polarization contrast (right) of standing water on the Drillfield.

In the image, we can see that the standing water has a high degree of polarization but a low polarization contrast. When we view the same standing water from different angles in Figure 65 and Figure 66, we can see how the degree of polarization changes.




Figure 65: S0 image (left), degree of polarization (center), and polarization contrast (right) of standing water on the drillfield.




Figure 66: S0 image (left), degree of polarization (center), and polarization contrast (right) of standing water on the drillfield.

The change in polarization of the standing water agrees with the results shown by Rankin and Matthies in 2008[24]. The apparent polarization of standing water is highly dependent on the angle the camera makes with the water and the position of the sun in the sky. The next set of images was taken at a higher angle from the Pylons that overlook the Drillfield on the Virginia Tech campus.




Figure 67: Standing water and mud on the drillfield when viewed from above.

Both the muddy areas and the standing water are highlighted when viewed from a higher angle. This could be because the viewing angle to the water is closer to Brewster's angle than when viewed from the ground. The muddy path on the top of the image is also very visible. This could be an indication that a polarimetric camera on a UGV should be mounted as high as possible to create a viewing angle closer to Brewster's angle.

4.2.2 Moving water

An important obstacle of interest for an unmanned ground vehicle is small bodies of water that could pose problems, such as streams or rivers. The camera was taken to Stroubles Creek on the Virginia Tech campus to evaluate its performance in detecting water features. The results showed that moving water was easily detected using the degree of polarization and the polarization contrast functions.

The first set of images was taken looking down the length of Stroubles Creek in the early afternoon with the sun behind a cloudbank. The water in Stroubles Creek was running and not standing water. The sun was roughly in front of the camera looking towards the southwest at an altitude of 61° in the sky. The camera was approximately 10 feet above the water on the bridge that crosses over Stroubles Creek next to West Campus Drive. The raw images are shown in Figure 68.




Figure 68: The four images from each camera looking down Stroubles Creek.(top left) Horizontally filter (top right) Vertically filter (bottom left) 45° filter. (bottom right) -45° filter.

The raw images show Stroubles Creek as it moves away from the camera. The tree in the top right is approximately 100 feet away from the camera. The camera exposures were equalized manually before taking data by pointing the camera straight down at the ground at a concrete sidewalk. The intensity level was set to provide a well exposed image in S_0 . The Stokes images calculated from the images in Figure 68 are shown below.




Figure 69: The Stokes vectors calculated from the images in Figure 68. From the top left going clockwise: S0, S1, S3, S2.

The degree of polarization image calculated from the Stokes images is shown in Figure 70.




Figure 70: The degree of polarization calculated from the Stokes images in Figure 69.

We can immediately see that the degree of polarization has highlighted the water's surface compared to the surrounding grass. The surface of the water has values in the 0.4 to 0.6 range while the grass has values below 0.1. This strongly suggests that the light reflecting off the surface of the water is at least partially polarized while most of the light from the surrounding land is not. To make sure that the data from Camera 3 is not providing false data due to its corrected exposure, we removed the S3 term from the degree of polarization equation. The resulting image is shown in Figure 71.




Figure 71: The degree of polarization image calculated without the S3 term.

With only S0, S1, and S2 terms, the degree of polarization of the image still shows that the water's surface is different than the surrounding land, although at a muted intensity. The water's surface has values in the 0.1 to 0.3 range, while the grass still is below 0.1.

The "Polarization Contrast" method proposed by Matthies [24] and detailed in Section 2.4.1 is applied to the same scene in Figure 72.




Figure 72: A polarization contrast image of Stroubles Creek.

The polarization contrast only uses the horizontal and vertical polarization images, which assumes that there is a difference in polarization between them to make its distinction. This proves to be correct, with the average value of the water's surface above 0.1 and the average value of the grass around 0.01. This method would prove useful if the intensity differences between four cameras using the beam splitter assembly could not be easily resolved. A

simplified version of the polarization camera could implement this technique with just two cameras, two filters, and one beam splitter.

The degree of polarization of the water appears to be resistant to whether the sun was visible or behind clouds, as shown in Figure 73.




Figure 73: The degree of polarization of Stroubles Creek with the sun behind clouds (left) and the sun visible (right).

The degree of polarization remains high on the surface of the water in both images compared to the surrounding grass in both bright sunlight and with the sun behind clouds. The intensity of the degree of polarization was lower behind clouds, but still significantly higher than the grass.

The degree of polarization also appeared to be independent of the location of the sun relative to the body of water and to the camera. The camera was moved to the side of Stroubles Creek, shown in Figure 74.



Figure 74: The polarimetric camera at Stroubles Creek.

While the previous were taken with the water, camera, and sun all relatively lined up with each other, the scene in Figure 75 was taken with the Sun to the left of the camera. In the raw images are in Figure 75, the camera is looking roughly north-northwest while the sun is in the southwestern part of the sky at an elevation of about 58° .




Figure 75: Stroubles Creek viewed perpendicular to the Sun, located to the left from this point of view.

From the raw images, the degree of polarization was calculated, shown in Figure 76.



Figure 76: The degree of polarization of Stroubles Creek viewed perpendicular to the sun.

From these images, we see that the degree of polarization of the water's surface remained higher than the surrounding grass when viewed at a different angle relative to the sun.

4.2.3 Man-made objects

A polarimetric camera can add additional information for computer systems in a UGV to make better decisions. One common example for UGVs is road detection.



Figure 77: A pair of image of West Campus Drive at Virginia Tech. (left) Horizontal polarized image. (right) Vertically polarized image.

To the untrained eye, the images in Figure 77 look like typical images of a road passing through campus. If we apply a simple polarization contrast technique from Equation (2.3) to the pair of images, we get the resulting image in Figure 78.



Figure 78: A polarization contrast image of West Campus Drive on Virginia Tech's campus.

As we can see, the road clearly stands out from the rest of the scene. This is due to the angle the road is making with the sun, producing some glare on its surface. The walking path to the right of the road also stands out from the grass it cuts through. The painted lines on the road are also clearly visible as different from the road surface. We can also see that the building behind the road stands out, including through one of the trees, because of the light filtering through the leafless branches.

Besides roadways, cars are another typical object of interest for many UGV platforms. Human drivers often employ polarized sunglasses to reduce glare and make it easier to see while driving on sunny days. For a computer system, the polarization of reflections on a vehicle could potentially provide useful information. A sample image of a vehicle is shown in Figure 79 below.



Figure 79: A set of polarized images of a typical van. (top left) Horizontally polarized. (top right) Vertically polarized. (bottom left) 45° polarization. (bottom right) -45° polarization.

The scene of the vehicle in Figure 79 was taken in mid-afternoon with the sun oriented towards the left side of the image. Using these images, the Stokes parameter images can be calculated, shown in Figure 80.



Figure 80: The stokes parameters calculated from the images in Figure 79. (top left) S_0 . (top right) S_1 . (bottom left) S_2 . (bottom right) S_3 .

As we can see, the windshield of the van shows up as polarized the Stokes images. This polarized light is the same effect that is filtered out by polarizing sunglasses that are popular with drivers. A similar effect is seen in the cars in Figure 81.



Figure 81: Raw images of cars on West Campus Drive at Virginia Tech.

The polarization contrast image calculated from the horizontal and vertical images are shown in Figure 82.



Figure 82: The polarization contrast image of the same scene shown in Figure 81.

Looking at the polarization contrast, we can clearly see that the windshields of the cars facing the camera are highlighted compared to the rest of the scene. The minivan in the center of the image, however, does not, indicating that this polarization could be dependent on the angle of the glass to the camera relative to the sun. However, this still is potentially useful information for a UGV looking to classify objects in a road environment.

5 Results

The experiments in Chapter 4 proved the feasibility of using a beam splitter system as the core of a low cost polarimetric camera. With a few preparation steps, the current camera could calculate real-time polarization data and display it to a user or save it to disk. The beam splitter design used in this thesis could become the core of a future camera design for use on a UGV. Chapter 5 will review the results from the experiments of the previous chapter.

5.1 Camera Alignment

The requirements in Section 3.2.1 stated that the images from each of the cameras in the overall polarimetric camera system had to be registered to each other so that the same pixel in each camera corresponded to the same point in space.

The alignment tests in section 4.1.1 showed that the image alignment error between the Firefly cameras reached a steady state once a certain distance from the camera was achieved. After the object was at least 25 feet away, the amount of pixel shift needed to align the images to each other in software stayed constant, whether the object was at 30 feet or 50 feet from the camera.

This would allow for a future polarimetric camera to have a preset minimum effective distance for it to function, similar to how a current camera can have its focus distance set beforehand. Once this minimum distance is established, calibration can be permanently set and remain valid for any distances beyond this minimum distance while the vehicle is used. This minimum distance is also related to the focal length used, which can be adjusted based on future requirements set by a new UGV development program. This required calibration is similar to the calibration that has been necessary with current camera sensors on the UGVs previously tested at the Mechatronics Lab at Virginia Tech.

5.2 Artificially Polarized Light Source

Once a method for camera alignment was developed, the next test took images of a light source with known polarization, an LCD monitor at -45° linear polarization. The Stokes vector of a fully polarized source can be easily calculated, as shown in Section 3.1.3. When images of the LCD monitor were taken with the polarimetric camera, it produced images that correlated with what was expected. The -45° filtered image let almost all the light from the monitor through, the 45° image stopped most of the light from passing through, while the 0° and 90° filtered images let about half as much light in as the -45° image. This set of images produced a set of Stokes images that correctly corresponded with a -45° polarized light source. The camera was then ready to begin collecting data from outdoor scenes.

5.3 Water and Mud

The data taken with the polarimetric camera showed potential for detecting mud and standing water surrounded by grass. The degree of polarization and polarization contrast of mud had a high amount of contrast when compared to the surrounding ground. For a UGV, this is potentially very useful information.

When the polarimetric camera observed standing puddles of water, results were mixed. The degree of polarization and polarization contrast of the water stood out from the surrounding grass. However, the intensity that the water stood out varied greatly depending on the position of the sun with respect to the camera and the water. This agrees with the results of Rankine and Matthies discussed in Section 2.4.1.

Images taken of running water proved to be different than standing water. The degree of polarization and polarization contrast images calculated from the Stokes images taken by the camera showed that the running water stood out very clearly from the surrounding ground. The degree of polarization and polarization contrast of the running water remained high even when looking at the water from different angles, as well as with the sun both behind clouds and shining directly. This indicates that the ability to detect running water could be more robust than standing water.

5.4 Man Made Objects

Finally, the polarimetric camera observed scenes that had roads and vehicles in the image. A special case was observed where an inclined road produced visible glare in the 0° polarization image but almost no glare in the 90° polarization image. When the polarization contrast of the scene was calculated, the road clearly stood out from the rest of the scene.

Data taken at other times of day showed that the degree of polarization of highly reflective objects such as car windshields also stood out from the rest of the scene. Polarizing sunglasses are commonly used by drivers to cut down on glare on the road. The polarimetric camera was able to detect this and display it in the form of a higher degree of polarization. This could be used on a UGV that travels in an urban environment to provide more information for the UGV about objects that it is detecting.

6 Conclusion

6.0 Summary of Chapter

Chapter 6 will summarize the methods used to design and assemble the polarimetric camera. It will then review the results from the polarization data collected by the camera. Finally, chapter 6 will review potential future work for the next generation of the low cost polarimetric camera and areas of improvement to the camera itself.

6.1 Polarimetric Camera Design and Assembly

Overall, the polarimetric camera design was successful. It provided real time polarimetric data at an order of magnitude less cost than commercially available cameras. It was able to use four different cameras to look at the same scene with different polarizing filters applied to each input. The beam splitter method of registering images worked with some drawbacks. The beam splitter alignment required careful machining of the beam splitter assembly with a relatively high degree of accuracy. The beam splitters also had to be secured using a two-part epoxy, making any errors permanent and any changes to their alignment impossible once set. The base plate had a $\frac{1}{4}$ -20 tapped hole in the center that allowed it to be easily mounted on a standard tripod for collecting data.

The major issue with the beam splitter assembly was the large change in intensity between cameras looking at different outputs of the assembly. This occurred because the beam splitters used in the camera are not ideal and reflect and transmit light in different amounts. Cameras 2 and 3, which light that was reflected in the first beam splitter, were noticeably darker than Cameras 0 and 1, which viewed light that was transmitted through the first beam splitter. Camera 3, which viewed light that was reflected twice, was the darkest out of all the cameras.

In order to be used as a sensor in a UGV, the intensity difference must be corrected, either physically with a different beam splitter design, or in software by characterizing the difference and applying a correction. For this prototype, manually correcting the camera intensities before taking data was sufficient and provided good polarimetric results. The camera intensities were corrected by observing a uniform, unpolarized surface, such as looking straight down at grass outside or at a blank wall indoors. The mean intensity and histogram of each camera feed was calculated and displayed to the user in real time to aid this process.

The filter holder system, once linear polarizing film was used instead of uncounted glass polarizers, worked well. It was difficult to accurately set tolerances for the additive manufacturing process used to create the filter box, which made setting glass polarizing filters difficult as some of the openings were slightly too small. To work around this problem, linear polarizing film was adhered to the outside of the opening instead forcing a glass filter into the opening as originally designed. The filter holder box also kept out stray light while protecting the beam splitter assembly from errant touches.

6.1.1 Overall Cost

The first requirement listed in Section 3.2.1 was that the polarimetric camera had to be low compared to commercially available cameras. A summary table of the component cost of the polarimetric camera can be seen in Table 5

Table 5: List of components in the polarimetric camera and their part cost

Component	Price	How many	Total
Point Grey Firefly MV Camera	\$250.00	4	\$1000.00
ThorLabs 50/50 20mm non-polarizing beam splitter	\$181.50	3	544.50
Cargille Refractive Index Matching Fluid	\$63.75	1	\$63.75
Thorlabs 2x2" linear polarizing film	\$8.00	4	\$32.00
Aluminum and Machine Shop	\$317.96	1	\$317.96
Right angle USB Cables	\$8.85	4	\$35.40
7-port USB Hub	\$28.99	1	\$28.99
TI MSP Launchpad	\$9.99	1	\$9.99
16mm CS-mount lens	\$13.95	4	\$55.80
Total Cost			\$2088.39

The largest portion of the total cost was the Firefly MV cameras, followed by the beam splitters and the cost of machining the aluminum base plate. The total component cost of the polarimetric camera was \$2088.39. This does not include the cost of the filter holder, which was printed on a Virginia Tech FDM machine. If it was purchased commercially, it would add approximately \$31 to the total cost. The cost of software is free as both the QT IDE and OpenCV libraries are freely available at no cost. The Point Grey FlyCapture SDK is free to download from the Point Grey website.

At an order of magnitude less than what was quoted for a commercially available polarimetric camera, the camera designed and build for this thesis achieved its goal of low cost by a large margin. Even with higher quality parts, a future version of the camera would still come in well below commercial cost and provide more flexibility in how it is used. If the camera was produced commercially, larger volume builds would further lower cost as price breaks come into effect when ordering higher quantities of components.

6.2 Future work

The polarimetric camera constructed for this thesis has proved to be a functional sensor design that could easily be adapted for use on a future unmanned ground vehicle platform. There are several areas where future work for the polarimetric camera system can be targeted, which are detailed below. The current four camera design has several improvements that could be made to improve its results. There is more work that can be done in obtaining larger data sets of filtered images of a wide variety of scenes and locations to continue to test the camera's effectiveness and to provide a large enough data set for future machine learning algorithms to use.

6.2.1 Redesign of beam splitter platform

In future iterations of the polarimetric camera, the beam splitter platform could be redesigned to take up less space. The current beam splitter platform design was originally sized to incorporate a mechanical clamp to hold the beam splitter assembly together. However, the clamping

mechanism did not work as designed, necessitating the use of epoxy instead. Epoxy holds the beam splitter assembly together very well and would be the primary method of attachment in a future iteration of the camera. The areas that would be eliminated from the beam splitter platform are located around the edges of the platform.

The largest areas that would be removed are on the left side of Figure 83, on the back side of the beam splitter assembly.




Figure 83: A CAD screenshot of the beam splitter platform and assembly.

The new beam splitter platform edge would be lined up with or only slightly beyond the back face of the beam splitters. On the front side of the beam splitter platform, towards the right of the figure, the raised stops that the beam splitters are aligned to could be reduced in thickness. Overall, several millimeters of space would be saved in each direction, which would directly translate in being able to position the cameras closer to the beam splitters themselves, potentially allowing for a wider field of view and reducing the overall space of the polarimetric camera system. It would also allow for larger lenses to be used.

6.2.2 Redesign of filter holder

The current filter holder used on the polarimetric camera worked fairly well. It held the polarizing filters and protected the beam splitter assembly both from inadvertent contact by a user as well as shielded it from ambient light. The filter holder was printed out of ABS plastic on a Stratasys 768 FDM at Virginia Tech and would be inexpensive to have made commercially. The filter holder fit snugly onto the beam splitter platform, requiring the use of no screws to hold it down and made it easy to install and remove.

In the future, some improvements can be made to make orienting linear polarizing filters easier, as well as allow filters to be easily replaced. Currently, the linear polarizing filters are fixed directly to the filter holder box and cannot be removed. A drop-in or screw-in filter would be a better solution for the next iteration of the design.

The beam splitter assembly, coupled with the filter holder, allows for the basic design of the camera to be used for other research, not just polarimetric. For example, because the Firefly MV cameras are sensitive to both visible and near-infrared light, bandpass filters could be placed in front of each camera to segment out different wavelengths of interest. It would also allow a user

to swap different sets of filters for different experiments or tests without the need for multiple filter boxes or any adhesives. The polarimetric camera would then become a multispectral camera. Changing the filter holder design to facilitate this would be a beneficial next step in a future camera.

6.2.3 Alternative camera and lens selection

While redesigning the beam splitter platform could allow for some space savings, the largest contributor to volume consumed was the Firefly MV cameras and lenses. In the future, a board level camera could be used, such as the Point Grey Dragonfly2, which provides similar sensing characteristics in a smaller form factor. The smaller size comes at the expense of some protection to the camera components, as well as more difficulty in mounting the camera itself.

The current lenses used on the polarimetric camera, while inexpensive, had some issues with image quality and shifting as the focus ring was turned. In the future, spending more on higher quality lenses would prove to be very beneficial. Alternatively, if space was as a premium, a camera such as the Dragonfly2 could also accept microlens mounts, which would save space with smaller lenses.

6.2.4 Polarization data sets and other polarizing sources

The polarization camera could be used to obtain large quantities of data that would be useful for a machine learning algorithm. Machine learning uses prior knowledge, such as a database of correctly classified images, to interpret and process new data [45]. A machine learning algorithm could use the polarization data from the camera as one of many sources of information for object classification, an important part of autonomy for unmanned systems. The new information provided by the Stokes images from the camera could provide the extra data needed to classify areas of an image that traditional camera systems could not classify on their own.

The polarization data sets could be expanded to cover subjects beyond simply the water and vehicles discussed in this thesis. Any shiny, reflective object has potential to be detected using a polarimetric camera. This could include windows in buildings, sunglasses on a person, certain paint finishes, and more. In a different direction from pure obstacle sensing, a future iteration of the polarimetric camera could be used as a navigation tool using the polarization of the sky [46], much like the honeybees discussed in Section 2.2.1.

6.2.5 Simplified design

Images of water in section 4.2.2 showed that detection of moving water could be done with just horizontal and vertical filters using polarization contrast. Section 4.2.3 showed the potential for detecting automobile windshields using the same method. A simpler polarimetric camera could be designed using only one beam splitter and two cameras. With two linear polarizing filters instead of four, the simpler design would not be able to fully characterize the state of polarization of a scene, but would have less sources of error with image intensity and camera alignment. A simpler design could also allow for the use of wider angle lenses, as they would only have to look through one beam splitter instead of two.

References

- [1] Bossa Nova Tech, “SALSA Evaluation (email).” 2012.
- [2] C. Reinholtz, “DARPA Grand Challenge,” 2004.
- [3] C. Reinholtz, B. Leedy, J. Putney, J. M. Webster, and C. Bauman, “Virginia Tech Grand Challenge Team; DARPA Grand Challenge 2005,” 2005.
- [4] B. M. Leedy, J. S. Putney, C. Bauman, S. Cacciola, J. Michael Webster, and C. F. Reinholtz, “Virginia Tech’s twin contenders: A comparative study of reactive and deliberative navigation,” *Journal of Field Robotics*, vol. 23, no. 9, pp. 709–727, Sep. 2006.
- [5] A. Bacha, C. Bauman, R. Faruque, M. Fleming, C. Terwelp, D. Hong, A. Wicks, T. Alberi, D. Anderson, S. Cacciola, P. Currier, A. Dalton, J. Farmer, J. Hurdus, S. Kimmel, P. King, A. Taylor, D. Van Covern, and M. Webster, “Odin : Team VictorTango’s Entry in the DARPA Urban Challenge,” *Journal of Field Robotics*, vol. 25, no. 8, pp. 467–492, 2008.
- [6] US Navy, “US Navy Marines Test Autonomous Unmanned Ground Vehicles,” *DefenseTalk*, 2010.
- [7] TORC Robotics, “Ground Unmanned Support Surrogate,” 2012. [Online]. Available: <http://www.torcrobotics.com/case-studies/ground-unmanned-support-surrogate>. [Accessed: 25-Mar-2013].
- [8] P. Morton, “Research,” *Australian Centre for Field Robotics*, 2013. [Online]. Available: <http://www-personal.acfr.usyd.edu.au/p.morton/research>. [Accessed: 15-Mar-2013].
- [9] S. D. Buluswar and B. a. Draper, “Color machine vision for autonomous vehicles,” *Engineering Applications of Artificial Intelligence*, vol. 11, no. 2, pp. 245–256, Apr. 1998.
- [10] M. Rosenbiuma and B. Gothardb, “Getting More from the Scene for Autonomous Navigation : UGV Demo III Program,” vol. 3838, no. September, pp. 176–187, 1999.
- [11] Joan Turville-Petre, *The story of Rauð and his sons (Trans.)*. Viking Society for Northern Research, 1947.
- [12] R. Hegedüs, S. Åkesson, R. Wehner, and G. Horváth, “Could Vikings have navigated under foggy and cloudy conditions by skylight polarization? On the atmospheric optical prerequisites of polarimetric Viking navigation under foggy and cloudy skies,” *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 463, no. 2080, pp. 1081–1095, Apr. 2007.

- [13] B. Kahr and K. Claborn, “The lives of Malus and his bicentennial law.,” *Chemphyschem : a European journal of chemical physics and physical chemistry*, vol. 9, no. 1, pp. 43–58, Jan. 2008.
- [14] H. E. W. Jenkins, Francis A., *Fundamentals of Optics*. 1976.
- [15] G. Horváth and D. Varju, *Polarized Light in Animal Vision: Polarization Patterns in Nature*. 2004.
- [16] G. Horváth, J. Gál, and R. Wehner, “Why Are Water-Seeking Insects Not Attracted by Mirages? The Polarization Pattern of Mirages,” *Naturwissenschaften*, vol. 84, no. 7, pp. 300–303, Jul. 1997.
- [17] R. Schwind, “Spectral regions in which aquatic insects see reflected polarized light,” *Journal of Comparative Physiology A*, vol. 177, no. 4, pp. 439–448, Oct. 1995.
- [18] N. Lefaudeux, N. Lechocinski, S. Breugnot, and P. Clemenceau, “Compact and robust linear Stokes polarization camera,” in *Proceedings of SPIE*, 2008, vol. 6972, p. 69720B–69720B–12.
- [19] L. B. Wolff, “Polarization camera for computer vision with a beam splitter,” *Journal of the Optical Society of America A*, vol. 11, no. 11, pp. 2935–2945, Nov. 1994.
- [20] FluxData, “FD-1665 3-CCD Multispectral Camera,” no. 800. 2010.
- [21] FluxData, “Imaging Polarimeters,” 2013. [Online]. Available: <http://www.fluxdata.com/imaging-polarimeters>. [Accessed: 10-Jan-2013].
- [22] FluxData, “FD-1665 (email).” p. 6583, 2012.
- [23] J. L. Pezzaniti, D. Chenault, M. Roche, J. Reinhardt, J. P. Pezzaniti, and H. Schultz, “Four Camera Complete Stokes Imaging Polarimeter,” in *Proceedings of SPIE*, 2008, vol. 6972, p. 69720J–69720J–12.
- [24] A. L. Rankin and L. H. Matthies, “Daytime mud detection for unmanned ground vehicle autonomous navigation,” 2008.
- [25] A. L. Rankin and L. H. Matthies, “Passive Sensor Evaluation for Unmanned Ground Vehicle,” *Journal of Field Robotics*, vol. 27, no. 4, pp. 473–490, 2010.
- [26] L. Matthies, P. Bellutta, and M. Mchenry, “Detecting water hazards for autonomous off-road navigation,” in *Proceedings of SPIE*, 2003, vol. 5083, no. 2, pp. 231–242.
- [27] A. L. Rankin, L. H. Matthies, and P. Bellutta, “Daytime water detection based on sky reflections,” in *2011 IEEE International Conference on Robotics and Automation*, 2011, pp. 5329–5336.

- [28] D. C. West, “Disturbance detection in snow using polarimetric imagery of the visible spectrum,” 2010.
- [29] Y. Y. Schechner, S. G. Narasimhan, and S. K. Nayar, “Polarization-based vision through haze.,” *Applied optics*, vol. 42, no. 3, pp. 511–25, Jan. 2003.
- [30] NOAA, “Remote Sensing,” *National Weather Service*, 2010. [Online]. Available: http://www.srh.noaa.gov/jetstream/remote/remote_intro.htm. [Accessed: 18-Jan-2013].
- [31] S. Ruzin, “Polarized light,” 2012. [Online]. Available: <http://microscopy.berkeley.edu/courses/tlm/plm/pol.html>. [Accessed: 15-Jan-2013].
- [32] Japan Atomic Energy Research Institute, “Development of Soft X-ray (1-keV Region) Polarizer Using a Mica Crystal,” 2005. [Online]. Available: http://jolisfukyu.tokai-sc.jaea.go.jp/fukyu/tayu/ACT05E/img/05IMG/05_11.jpg. [Accessed: 15-Jan-2013].
- [33] E. Hecht, *Optics*. 2002, pp. 374–379.
- [34] Wikipedia, “Brewster’s angle,” *Wikimedia Foundation, Inc.*, 2007. [Online]. Available: <http://en.wikiversity.org/wiki/File:BrewsterAngle.jpg>. [Accessed: 16-Jan-2013].
- [35] A. Duffy, “Polarization by Reflection,” 2005. [Online]. Available: http://webphysics.davidson.edu/physlet_resources/bu_semester2/c27_brewster.html. [Accessed: 17-Jan-2013].
- [36] Wikipedia, “Water Reflectivity,” 2010. [Online]. Available: http://en.wikipedia.org/wiki/File:Water_reflectivity.jpg. [Accessed: 05-Jul-2013].
- [37] “One-Way Glass Stops Glare,” *Popular Mechanics*, pp. 481–483, Apr-1936.
- [38] OpenStax College, “Polarization,” 2012. [Online]. Available: <http://cnx.org/content/m42522/latest/>. [Accessed: 18-Jan-2013].
- [39] Edmond Optics, “Broadband Non-Polarizing Cube Beamsplitters,” 2013. [Online]. Available: <http://www.edmundoptics.com/optics/beamsplitters/cube-beamsplitters/broadband-non-polarizing-cube-beamsplitters/2239>. [Accessed: 12-Mar-2013].
- [40] D. G. Lowe, “Object recognition from local scale-invariant features,” *Proceedings of the Seventh IEEE International Conference on Computer Vision*, pp. 1150–1157 vol.2, 1999.
- [41] VLfeat.org, “Tutorials - SIFT,” 2013. [Online]. Available: <http://www.vlfeat.org/overview/sift.html>. [Accessed: 12-Mar-2013].
- [42] ThorLabs, “BS016 - 50:50 Non-Polarizing Beamsplitter Cube, 400 - 700 nm, 20 mm,” 2013. [Online]. Available:

- <http://www.thorlabs.com/thorProduct.cfm?partNumber=BS016>. [Accessed: 13-Mar-2013].
- [43] I. Point Grey Research, “FireFly@MV CMOS Camera,” 2012. [Online]. Available: http://www.ptgrey.com/products/fireflymv/fireflymv_usb_firewire_cmos_camera.asp.
- [44] Texas Instruments, “MSP430 LaunchPad Value Line Development kit,” 2013. [Online]. Available: <http://www.ti.com/tool/msp-exp430g2>. [Accessed: 13-May-2013].
- [45] D. Barber, *Bayesian Reasoning and Machine Learning*. Cambridge University Press, 2012.
- [46] J. a North and M. J. Duggin, “Stokes vector imaging of the polarized sky-dome.,” *Applied optics*, vol. 36, no. 3, pp. 723–30, Jan. 1997.

APPENDIX A: Typical Stokes and Jones vectors for polarized light

Polarization State	Stokes vectors	Jones Vectors
Horizontal polarization	$\begin{bmatrix} 1 \\ 1 \\ 0 \\ 0 \end{bmatrix}$	$\begin{bmatrix} 1 \\ 0 \end{bmatrix}$
Vertical polarization	$\begin{bmatrix} 1 \\ -1 \\ 0 \\ 0 \end{bmatrix}$	$\begin{bmatrix} 0 \\ 1 \end{bmatrix}$
45° polarization	$\begin{bmatrix} 1 \\ 0 \\ 1 \\ 0 \end{bmatrix}$	$\frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \end{bmatrix}$
-45° polarization	$\begin{bmatrix} 1 \\ 0 \\ -1 \\ 0 \end{bmatrix}$	$\frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ -1 \end{bmatrix}$
Right-handed circular polarization	$\begin{bmatrix} 1 \\ 0 \\ 0 \\ 1 \end{bmatrix}$	$\frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ -i \end{bmatrix}$
Left-handed circular polarization	$\begin{bmatrix} 1 \\ 0 \\ 0 \\ 1 \end{bmatrix}$	$\frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ i \end{bmatrix}$

APPENDIX B: Typical Jones and Mueller Matrices

Linear Optical Element	Jones Matrix	Mueller Matrix
Horizontal Linear Polarizer	$\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$	$\frac{1}{2} \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$
Vertical Linear Polarizer	$\begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}$	$\frac{1}{2} \begin{bmatrix} 1 & -1 & 0 & 0 \\ -1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$
45° Linear Polarizer	$\frac{1}{2} \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}$	$\frac{1}{2} \begin{bmatrix} 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$
-45° Linear Polarizer	$\frac{1}{2} \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}$	$\frac{1}{2} \begin{bmatrix} 1 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 \\ -1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$
Quarter-wave plate, fast axis vertical	$e^{i\pi/4} \begin{bmatrix} 1 & 1 \\ 1 & -i \end{bmatrix}$	$\frac{1}{2} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & -1 & 0 \end{bmatrix}$
Quarter-wave plate, fast axis horizontal	$e^{i\pi/4} \begin{bmatrix} 1 & 1 \\ 1 & i \end{bmatrix}$	$\frac{1}{2} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & -1 & 0 \end{bmatrix}$
Homogeneous right-handed circular polarizer	$\frac{1}{2} \begin{bmatrix} 1 & i \\ -i & 1 \end{bmatrix}$	$\frac{1}{2} \begin{bmatrix} 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 1 \end{bmatrix}$
Homogeneous left-handed circular polarizer	$\frac{1}{2} \begin{bmatrix} 1 & -i \\ i & 1 \end{bmatrix}$	$\frac{1}{2} \begin{bmatrix} 1 & 0 & 0 & -1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ -1 & 0 & 0 & 1 \end{bmatrix}$

APPENDIX C: MSP430 Source Code

The MSP430 trigger board was programmed in C using Texas Instruments' Code Composer Studio.

```
/*
 * Simple Camera Trigger - Mark Umansky - Spring 2013 *
 * Notes: 1Mhz clock
 */

#include <msp430.h>
int main(void) {
    WDTCTL = WDTPW + WDTHOLD;           // Stop watchdog timer
    P1DIR = 255;                        // Set P1.0 to output direction
    P1OUT = 0;
    while(1) {
        P1OUT = 255;
        __delay_cycles(5000);
        P1OUT = 0;
        __delay_cycles(25000); }
}
```

APPENDIX D: Camera capture source code

The source code used to control the polarimetric camera and process the data was developed in C++ using the QT Project IDE, which has an easy to use GUI creator. It requires the Point Grey FlyCapture2 SDK, QT 5.0.0 IDE, and OpenCV 2.4.3.

The source code for the software written for this thesis can be found at:
<http://filebox.vt.edu/users/mumansky/thesis/>

APPENDIX E: Procedure for using the polarimetric camera

There are several short steps needed to set up the prototype polarimetric camera before data collection can begin. The entire process can be completed in a few minutes.

1. Power on laptop, plug in USB cable.

The polarimetric camera receives both data and power through the single USB output from the onboard USB hub. Connecting the laptop to the USB hub automatically powers on and starts the cameras. The trigger from the MSP430 will also automatically start sending signals to the cameras.

2. Open control software and start the program.

The control software will automatically find the cameras and start receiving data. If a camera is not connected, the software will prompt you to check all cables and exit.

3. Focus each lens until subject is in focus across all cameras.

This step is required to ensure that the scene is in focus across all cameras, providing the best possible image quality. The raw images displayed in the software can be used to check focus.

4. Perform image alignment.

The image alignment discussed in Section 4.1.1 is required to provide accurate data for the Stokes vector calculation. The image alignment is performed by looking at the subtracted images in the tabs highlighted by the red box in Figure 84 and adjusting the correction for Cameras 1 through 3 by comparing them to Camera 0. The alignment controls are highlighted by the blue box.

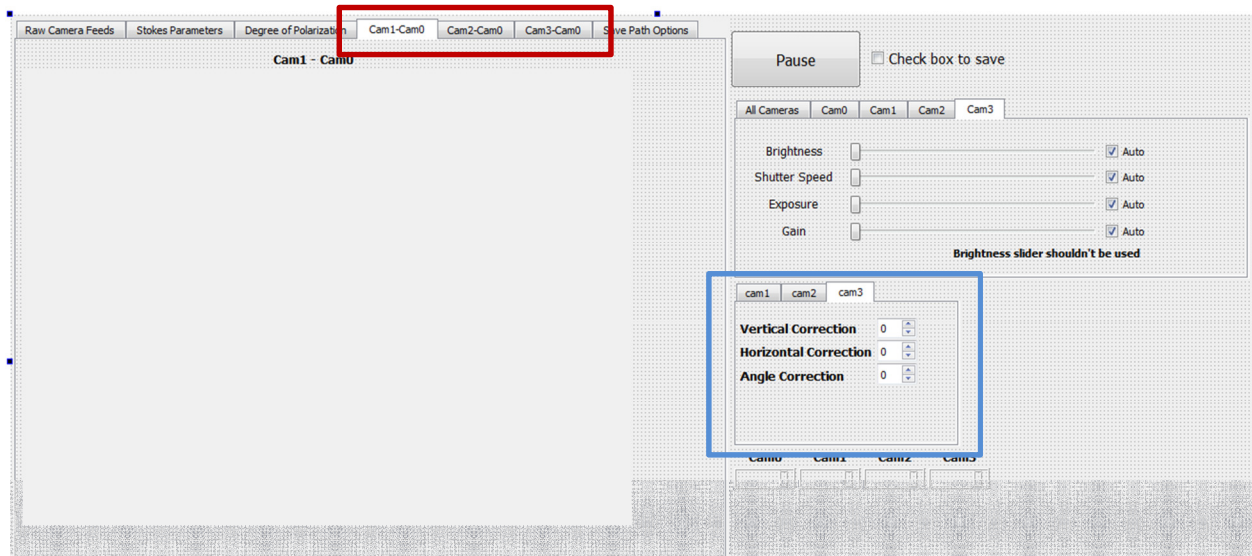


Figure 84: A screenshot of the control software with image alignment panels highlighted in blue and alignment image tabs highlighted in red.

For a future camera mounted on a UGV, step 4 would only need to be performed once. Once alignment would be set, it would be valid for any objects at its minimum focus distance and beyond.

5. Adjust exposures on non-polarized light source.

When the cameras first power on, they will automatically expose by default. In the raw images, their exposures will look the same. In reality, the auto exposure algorithm is adjusting each camera's parameters individually because of the non-ideal beam splitter assembly, which destroys any intensity differences due to polarization. Thus, manually setting each camera's exposure is necessary before taking data. The camera parameters for setting exposure are highlighted by the blue box in Figure 85.

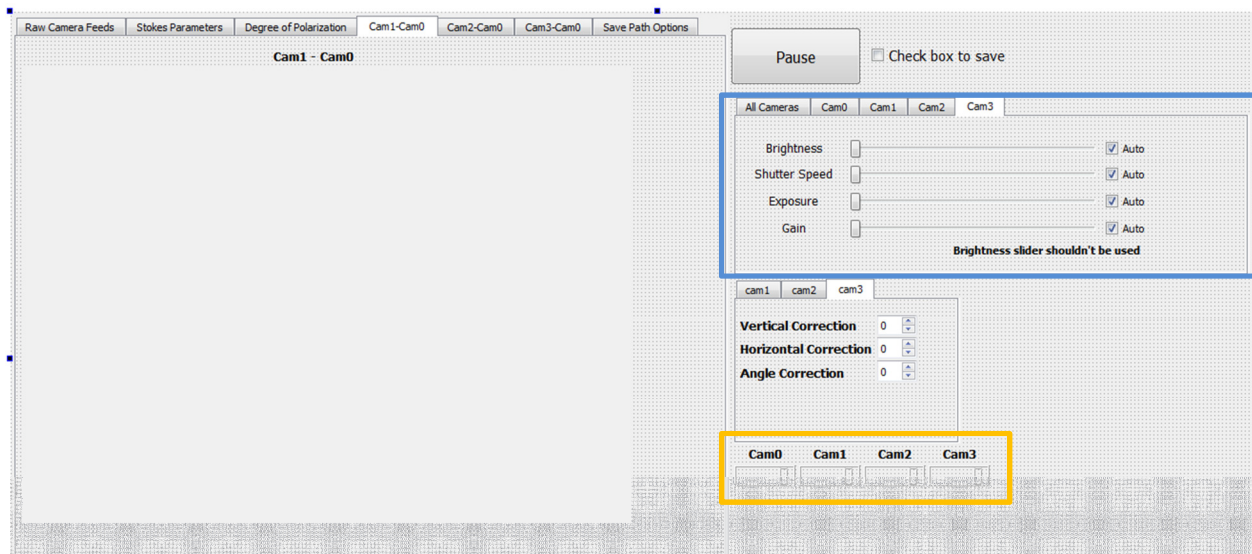


Figure 85: A screenshot of the control software with camera exposure controls highlighted by the blue box.

To adjust exposure, point the camera at a non-polarized light source, such as straight down looking at grass or concrete. Adjust each camera's exposure parameters for shutter speed, exposure, and gain until the average intensities from each image are as close as possible. The average intensities are highlighted by the orange box in Figure 85. Once equalized, the camera is ready to be pointed up at the scene of interest. In the future, a beam splitter assembly that more closely splits light 50:50 would reduce or eliminate the need for this step.

6. Save data to disk.

Once the camera is pointed at the scene, it is ready to save data to disk. The save path must be set beforehand so the program knows what folder to save images in. Checking off the save box in the top right corner will start saving all images in real time until the box is unchecked.