Purdue Aerial Robotics 2018

Journal Paper for AUVSI Student UAS Competition

Abstract

This paper describes the Purdue Aerial Robotics UAS for use in the 2018 AUVSI SUAS competition. As a member committee of Purdue's IEEE student branch, Aerial Robotics has a heavy focus on low level design, with a student-built fixed wing aircraft and custom software solutions. The aircraft itself is designed to be sturdy and easily controllable, while the Pixhawk flight controller provides robust control of the plane's flight path. The team uses Mission Planner, a commercially available ground control station (GCS), to control the aircraft from a laptop on the ground. The video processing system is a machine learning based target recognition system run almost entirely from an NVIDIA Jetson TK1 mounted on the UAS. It has the ability to correctly identify shapes, characters, and colors by name along with an object's relative image location and orientation. For the duration of the season, safety and reliability were the team's primary concerns, both in development and in operation of the UAS.

1. **Systems Engineering Approach**

The development of the UAS began with understanding the subsystem requirements for each competition task. Electrical and hardware requirements were developed to satisfy the basic competition tasks such as powering and controlling the motors, reading and regulating the voltage and power from the system, and selecting the proper camera to view the distinct targets on the ground. The final hardware design is incorporated into the airframe design as the required payload. The airframe is designed to hold the electronics in place and ensure the safety of critical components using a stable frame configuration and to have a modular interface for easy access to component testing and replacement. The software specifications were listed and developed in order of priority for each competition task and requirement using the Agile/Scrum methodology. Each component in the design was tested, debugged, and verified to meet the specifications determined at the beginning of the design phase.

Three teams were created to develop the main UAS components: electrical, software, and aero-mechanical. Once the requirements for the overall system were agreed upon, the individual teams would develop their subsystems with communication between teams to discuss or integrate components. The electrical and software teams would use a commercial airframe to test their subsystem components while the aero-mechanical team built and tested the constructed airframe. Final system integration and verification involved repeated flight testing to test key components in the design.

1.1 Mission Requirement Analysis

The mission situation as given by the competition description is as follows: "There is a lost hiker located in a remote area who has called for help. Search and rescue has tasked an Unmanned Aerial System (UAS) to find the hiker, direct rescue personnel to the hiker’s location, and deliver urgent aid supplies." For this mission, it is required that the airplane be able to be set up and used quickly, 20 minutes of setup and 40 minutes of flight, and cleaned up in only 10 minutes. The airplane itself must be under 55 lbs and fly no faster than 70KIAS at any time. Crucial to the competition is that the airplane be able to fly autonomously. It must be able to fly autonomously for 3 minutes minimum, and they encourage autonomous flight for the entire flight. Autonomous image recognition and hiker localization is optional but encouraged. This would include object detection and classification, which will be performed onboard the UAS. The aircraft, is required to remain within the bounds of the competition airport, including during its autonomous flight. It is also tasked with flying through waypoints and avoiding virtual obstacles autonomously. Of course, the aircraft is also required to fly safely along the way and without crashing.

Table 1. Target tasks, the percentage of possible total points, and expected points awarded.

Task	Percentage	Expected
Autonomous Flight	0.12	0.12
Waypoint Capture	0.03	0.03
Waypoint Accuracy	0.15	0.15
Stationary Obstacle Avoidance	0.1	0.0
Moving Obstacle Avoidance	0.1	0.0
Search Area	0.12	0.04
Off Axis Target	0.02	0.01
Interoperability	0.06	0.06
Air Delivery	0.1	0.0

Legend
Will Perform
Will Attempt
Will Not Attempt

1.2 Design Rationale

Continuing from the previous year's work, a list of tasks necessary for autonomous flight was written and a critical path was found to ensure that the flight controller would be finished by the middle of March. The electrical and software teams chose to initially base their work off of the previous years' custom PCB and autopilot software. They chose to pursue a custom design to encourage members to learn more about the low-level development process. By mid-year the team had decided to pursue a safer and more reliable software solution by using a Pixhawk for all flight control and automation requirements. This ultimately led to a higher-level design which allowed for faster development at the expense of lower-level development experience in embedded systems.

The software team needed to choose a camera to support their object detection requirements. The team needed to accurately detect objects with a required resolution of 1 pixel per inch from an altitude of 35 meters. The Z Camera E1 was chosen for its high resolution, with a 14mm lens. A neural network was chosen to classify images taken by the camera.

The aero-mechanical sub team sought to radically change from the previous plane to simplify the design with information gleaned from the last years build. At the start of the season there was an influx of new members who had to be taught composite manufacturing and SOLIDWORKS before they could help with the aircraft. As a result, much of the initial design work was on the shoulders of returning members. Due to a majority of the budget going to competition costs, the budget for the airframe was approximately $2,000 and utilized many resources from existing Purdue labs and classes. The purpose of the airframe was to be a stable platform for the electrical subsystems, so the focus of the design was on stability and loading. As a result, the plane was designed with a large wing, high wing, and large control surfaces.

2. System Design,

The Purdue Aerial Robotics 2018 UAS is a simplified version of the 2017 model, the details for which can be found in the 2017 Purdue Aerial Robotics AUVSI SUAS Journal Paper.

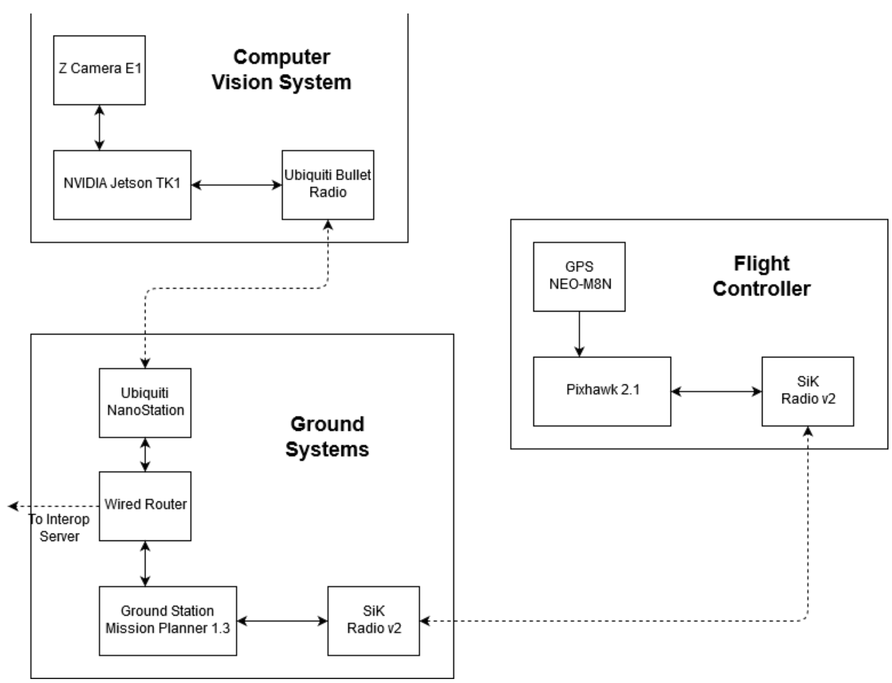

Figure 1. A diagram of the Purdue Aerial Robotics 2018 system.

2.1 Aircraft

The airframe is composed of a main wing assembly, tail assembly, and combined motor and electronics pod. The plane is designed to have a traditional tail configuration with a single motor and a rectangular wing. The main wing is a single section is made of a fiberglass skin around a foam core with a Kevlar reinforcement on the leading edge. In case of crashes, such a plane would be easier and faster to rebuild. Foam is also lightweight and easy to modify if changes were needed to be made to the plane. The wing acts as the main lift surface and structural component to house the pod and channel wiring to the servos, batteries, and tail. The tail was constructed in a similar fashion to the wing and was connected using a foam joint. The front landing gear is made of a stitch weave carbon fiber piece for strength and ease of acquisition and has two rubber wheels. A carbon fiber square tube is clamped to the motor mount and runs all the way to the tail with mountings for the main wing, electronics, and motor pod. Structural components such as the wing span and landing gears were designed with a factor of safety of 2 using a 3 g loading scenario.

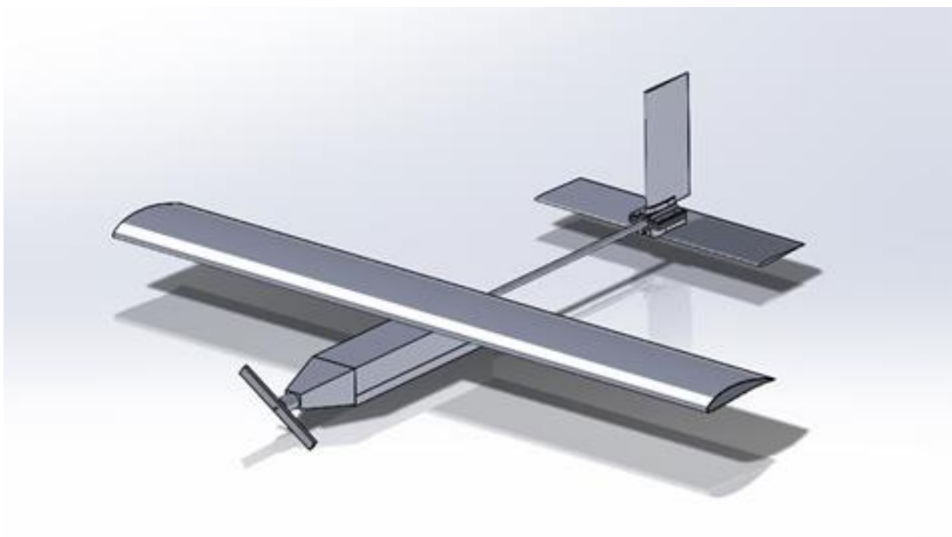

Figure 2. The isometric view of the main components and features of the Purdue Aerial Robotics UAS.

2.1.1 Fuselage

The motor and electronics pod acts as the fuselage and was fabricated from stitch weave fiber glass formed into a rectangular shape. The motor and electronics pod houses a motor, ESC, battery, propeller, Pixhawk, Jetson TK1, camera, and other essential electronics. The top section of the fuselage and two side walls were attached using epoxy, while the bottom section remained separable to allow for accessibility to the electronics of the aircraft. The bottom section was attached using Velcro strips as well as tape to ensure that it cannot fall off in flight. A tapered nose was added to the front of the fuselage to reduce drag and serve as a mounting surface for the motor. A slot was also cut out of the nose so that the battery could slide in, allowing the center of gravity to be shifted as needed.

2.1.2 Wing and Tail

To begin the design process, the total aircraft weight was estimated to be 5 kg after analyzing previous competition winners, who used designs comprised of composites with carbon fiber booms along with a straight wing configuration. Using this weight along with the stall speed requirement and the characteristics of the selected airfoil (MH38, chosen for the high lift-to-drag ratio and gentle stall) produced an estimated main wing area of approximately 0.6 m^2. A wingspan of 2.2 meters was chosen in order to maximize aspect ratio while remaining within structural and practical constraints.

XFLR5 analysis was then used to refine the design. A hot wire cutter was used to cut a MH38 airfoil with a chord of 0.27 m out of a Formula 250 R10 foam block. The airfoil shape is constant throughout since this greatly increases manufacturability with a foam wire cutter without substantial performance losses. A fiberglass layup was used to strengthen the wing. The ailerons were cut using a bandsaw and were connected at the hinge with tape. The servos were placed in slots that were milled out with a drill press to allow the servo to sit flush to the surface of the wing. Using constant-lift analysis, a cruise speed of 14 m/s was found to maximize the lift-to-drag ratio of the wing. The cruise lift-to-drag ratio given by XFLR5 VLM was found to be 25.

The horizontal and vertical tails were sized using tail volume coefficients. The values for these coefficients were selected based on the performance and handling of the previous year's aircraft. The tail was then included in the XFLR5 analysis to obtain an incidence angle for the horizontal tail and to verify that static stability had been achieved. A simple dynamic stability analysis was also conducted to ensure that the aircraft was safe to fly.

Table 2. Airframe properties.

Performance	
Aircraft Mass	5 kg
Cruise Speed	14 m/s
Stall Speed	9.92 m/s

Table 3. Wing and tail airfoil properties.

	Main Wing	**Horizontal Tail**	**Vertical Tail**
Area	0.594 m^2	0.116 m^2	0.05 m^2
Span	2.2 m	0.68 m	0.31 m
MAC	0.27 m	0.17	0.16
Aspect Ratio	8.15	3.99	1.92
Airfoil	MH38	NACA 0012	NACA 0012
Loading	82.37 Pa	-	-

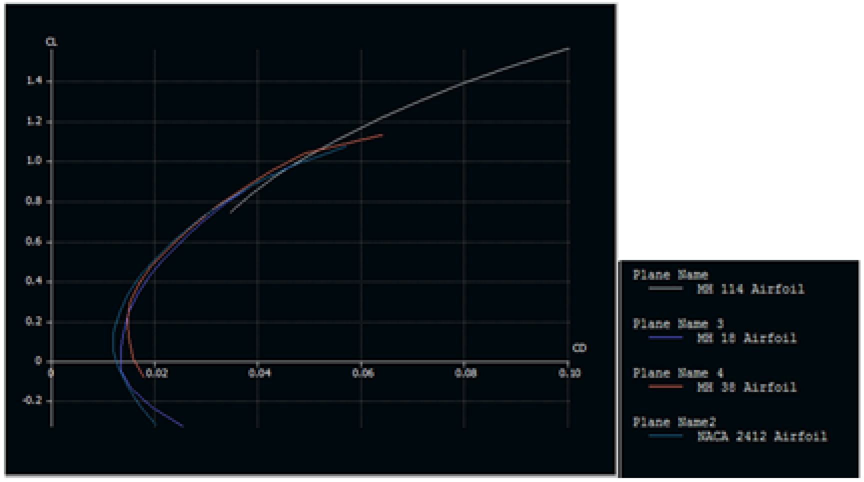

Figure 3. CL/CD plot from XFLR5

2.1.3 Propulsion

The motors are Sunnysky x2820 920kV coupled with a 12x6E APC propeller, powered by Multistar 4S 14.8V 16000mAh lithium-polymer battery and Turnigy brushless 80-amp ESC. The total required power output for cruise is 266W. This system is capable of providing 1510W.

2.2 Autopilot

The autopilot system currently uses a Pixhawk 2.1 for autonomous navigation. The Pixhawk system was chosen for its interoperability and ease of development. A GPS module, battery monitor, and the telemetry radio were easily integrated into the Pixhawk. The Pixhawk runs ArduPlane 3.8, chosen for its proven effectiveness and off-the-shelf capability. Using a tried-and-tested platform grants additional safety and reliability at no extra cost to development time.

The Ground Control Station (GCS) used is ArduPilot's Mission Planner. The team chose this GCS for the ease of integration with the Pixhawk.

Figure 4. Mission Planner (GCS)

2.2.1 Original Custom Autopilot

The team originally intended to develop a custom autopilot system. The purpose was to give students at Purdue the chance to develop lower-level code and algorithms to give them a greater depth of experience, especially with regard to the electrical and computer engineering students. Due to time constraints the team ultimately chose to fall back on the more reliable Pixhawk. However, the plans for the custom flight control system are

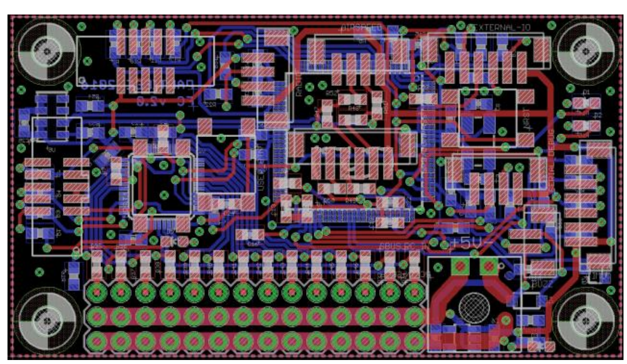

Figure 5. Original Custom PCB

documented here.

The autopilot system was originally built on a custom printed circuit board (PCB). It used an STM32F4 microcontroller to run the main custom control stack, as well as an STM32F1, used as a co-processor for manual takeover. The requirements for autonomous flight as described in the rules drove choices on design for the board. The flight controller used a 9-axis IMU, LIDAR rangefinder, GPS, barometer, and an airspeed sensor to collect information about the environment. The custom PCB was created to minimize the complexity of the wiring in the UAS by providing a unified access point for connecting sensors. This allowed for much greater freedom in implementing parallel interfaces such as those used by the main and coprocessor.

The main processor originally contained the flight controller, avionics, and ground control communication. The coprocessor acted as a “smart multiplexer” that passes either the outputs of the flight controller or the safety override controls to the control surfaces and propellers. The decision to separate the multiplexing of the signals into a separate processor was made so that if for any reason the main processor were to freeze or reset itself the plane could still be recovered through manual flight over the much simpler coprocessor. In normal operation, the flight controller signals were passed through to the control surfaces. However, when a switch was flipped on the safety pilot’s controller, the plane switched to manual control.

The flight control system was built on a STM32F4 microcontroller and featured a hand-built kernel written in C with support for interrupts and task scheduling. A list and breakdown of the tasks on the primary controller can be seen in Table 4. The custom flight control system allowed for team members to write custom functions that would not normally be available in commercial products that are specific to the team requirements.

Table 4. The tasks on the primary controller.

Task	Description
Telemetry Radio	Sends and receives communications between GCS and the Flight Controller
Serial Debug	Allows for debugging through the main controller’s serial port
User Button	Checks if the user button has been pressed
Sensors	Updates sensor information
Inter-Micro Communications	Sends control signals to the coprocessor
Flight PID	Control surface algorithm
Navigation	Calculates error values to feed into control algorithm

The horizontal navigation system utilized a lookahead control along with linear parameterization of the plane's path to send a roll error term to the control unit. The lookahead control system found the heading that would direct the plane towards a point 5 meters ahead of the point on the line defined by the UAS's cross-track-error. The roll error was then fed to the horizontal control algorithm, the Simulink representation of which is shown in Fig. 6.

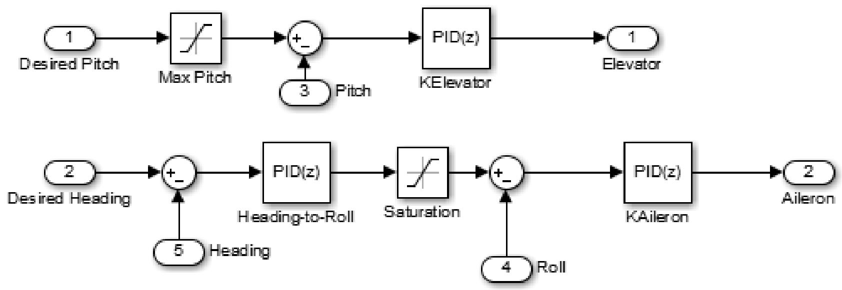

Figure 6: A Simulink representation of the horizontal control algorithm

Each of the PID blocks contains a proportional control. Due to the physical forces on the wings, most oscillations are naturally dampened. The saturation block for heading-to-roll prevented the aircraft from rolling more than 30 degrees in a turn. The max pitch clamp was set for 45 degrees to prevent the plane from stalling itself if given a waypoint that would require climbing at an extreme angle.

The gains for this algorithm were determined by developing a control algorithm in Simulink and testing it using a UDP interface with X-Plane 9. The stock model plane aircraft was used for sanity checking new developments to ensure that they would result in a stable controller given tuned constants. This prevented the team from pursuing non-viable controller designs and significantly sped up development of new control algorithms as the UAS did not have to be prepared to test new experimental features.

Due to a lack of time and a shortage of member involvement, the team decided to use the Pixhawk rather than a custom-built autopilot. This decision was made after weighing the tradeoffs that were explained above.

2.2.2 GPS

The GPS used is a NEO-M8N. The Pixhawk communicates with the GPS via the UBX protocol over a serial interface. The GPS provides the plane's current location as well as information about its height and ground speed that is fused with pressure and airspeed data. To prevent interference from other electronics, the GPS module is mounted in an isolated position on the outside of the plane.

2.3 Obstacle Avoidance

Obstacle avoidance is not yet implemented. However, the system has features to ensure that the UAS stays on course and does not leave the mission area. It will not allow a waypoint to be set to an area deemed to be off limits.

2.4 Imaging System

The Z Cam E1 is a 16-megapixel camera with a maximum resolution of 4640 x 3480 pixels. This allows 1-inch thick lettering to be resolved from over 150 feet in altitude. Since the UAS will be flying below this altitude, this camera is sufficient for object detection and classification.

With an average cruise speed of 14 meters per second, the camera needs to take an image at least once every 2.7 seconds. The Z Cam E1 satisfies this condition, having the ability to take full resolution pictures at 15 frames per second at 16-megapixel resolution.

The gimbal tilts the camera to keep it looking down as the airplane tilts. The gimbal assembly is 203mm long, 115mm tall, and 50mm tall. With the camera installed, the total height is 120mm. The gimbal can rotate the camera 41 degrees in the pitch axis and 24 degrees in the roll axis. The gimbal is driven by two HS-82MG servos, one for each axis of rotation, and uses bearings in each of the joints.

The roll axis is controlled by direct drive to minimize slack, but the pitch axis is controlled by a wire linkage because a direct drive would have been too wide to fit in the airplane. The center position in the pitch axis can be adjusted with an A-bend in the wire linkage. The camera is attached with a clamp that holds it at the base of the lens, which allows the autofocus dial to rotate freely. The gimbal assembly is mounted within the airplane's fuselage with Velcro, which allows its vertical position to be adjusted and makes it easy to remove.

2.5 Object Detection, Classification, Localization

The UAS detects, classifies, and localizes objects through four major processes. The system transfers images from the camera into memory, detects regions of interest (objects) on the ground, classifies various aspects of the object through the use of a neural network, and synthesizes a final output. The OpenCV library is used throughout the project for its wide range of highly optimized image processing algorithms. Additionally, the Keras Python module is used to build the neural network.

2.5.1 Image Capture

Images are first captured using the Z Cam E1. They are then transferred to the Jetson TK1 board via a USB connection and loaded into the main image classification program. These images are loaded using OpenCV.

2.5.2 Region of Interest Detection

In order to detect objects, regions must first be parsed out of the full-size image for classification. After researching various region proposal algorithms, it was decided that BING++ would be the best option. BING stands for Binarized Normed Gradients, and it was developed by The University of Oxford and Boston University in 2014. BING++ is a more optimized version of the original algorithm and was updated in 2017. This algorithm performs high quality region proposal at a very high speed. Processing images at approximately 300 frames per second, BING++ parses incoming images for regions of interest much faster than any single-shot machine learning approach can. Since the entirety of the Computer Vision System's processing occurs on the UAS itself, the speedup enabled by BING++ is essential.

2.5.3 Image Classification

As previously mentioned, image classification is performed using a neural network. Initial inspiration for the network architecture came from PVANet, a powerful yet efficient object detection network. PVANet is designed to achieve accuracy comparable to the largest, most accurate object detection neural networks while only using 10% of the processing power. Since the goal of this Computer Vision System is to perform all of the image processing on the plane, many parts of PVANet were replicated for efficiency.

Moreover, the team's machine learning solution improves on the efficiency of existing object detection networks. Instead of having multiple networks each detecting a different feature of a given input image, like shape, character, orientation, and color for example, the network reuses the information gained from lower neuron layers to identify many features. With one input image, the neural network outputs all data required by the competition. Reusing layers increases total network efficiency over multiple networks of similar accuracy by nearly 60%.

2.5.4 Synthesis

Due to the nature of the neural network outputting all required data, synthesis and data transfer to the ground station is fairly straightforward. All data is converted from the vector output provided by the neural network directly into the format needed for transfer. The data is then sent over a socket connection to the Ground Control Station.

2.6 Communications

The UAS has two main communications subsystems: the telemetry interface and imaging interface. The interfaces are separated to minimize the risk of losing telemetry from the UAS. The imaging interface uses a Ubiquiti Bullet HP M transceiver on the UAS, and a Ubiquiti Nanostation M on the ground. These transceivers operate using Wi-Fi and transmit computer vision results to the ground.

The telemetry interface uses two SiK v2 transceivers to communicate between the ground control and UAS. The communication interface is MAVLINK. The MAVLINK protocol was chosen because it is a commonly used industry standard and allows for the flight controller to interact with a variety of possible ground station systems. This lends the system some modularity as the flight controller and GCS can be swapped at any time. In addition, it cuts the cost in development time and increases safety by using proven communications protocols. There is also a manual control interface that is used when a safety pilot is needed.

The SiK v2 transmitters 433 MHz communication, and the imaging system uses 2.4 GHz Wi-Fi. The RC safety control uses 2.4 GHz FHSS communication.

2.7 Cyber-Security

As encryption was not critical to the functioning of the autopilot, it was not a priority and was not implemented. If outside changes to the flightpath were to occur, operators at the ground control station would likely take over with manual control. The manual RC control is as vulnerable as other RC control systems. However, since it is a frequency hopping spread spectrum interface, it is difficult to jam the controls or spoof messages without finding the frequency sequence used to control the manual flight algorithm.

3. Test and Evaluation Plan

3.1 Developmental Testing

The developmental testing plan for the UAS is designed to provide tasks of increasing complexity while minimizing risks to personnel or property. The first round of tests was conducted to discern if the preliminary aircraft design was capable of flight. The camera was not attached to the aircraft for this test. The flight lasted approximately two minutes and confirmed that the prototype airframe was capable of flight. The next round of flight tests included both the camera and the Pixhawk and was used to gather flight data from the Pixhawk and test the autonomous system. During the flight test, there were some issues with the autonomous flight as the waypoints were not set correctly in the program. These flight settings were cleared and redone for future flights. The data collected from the Pixhawk was analyzed for further improvements. Further flight tests continued over several months to improve the UAS airframe, software, and electrical systems. The data from these tests were subsequently logged and analyzed.

3.2 Individual Component Testing

3.2.1 Sensors

The GPS was checked outside using a google maps comparison. The location that the GPS reported was within 5 meters of the expected position.

3.2.2 GCS-Plane Communication

The plane and GCS communication has been run through stability tests to ensure their proper operation. The communication has also been range tested, and is able to maintain communications for every distance measured, the furthest of which was nearly 0.5 km.

3.2.3 Image Processing

Figure 7. Determining the shape, character, and color names from an example test image.

Above in Fig. 7, the image processing algorithms correctly identify the two shape objects with their correct shape name, character name, and color names. The orange jacket on the ground, however, is classified as an “Unknown Object” meaning that it is not a standard object.

3.3 Mission Testing Plan

Purdue Aerial Robotics’ mission plan is focused on validating the autopilot in a competition setting while using the image processing system to find targets. The system is to be run in conditions as close to those in the competition as possible given crew availability and flight conditions. The mission plan for every competition level test is decided on the basis of maximizing the quantity of tasks that can be completed on a given run so as to ascertain the team’s readiness to operate within various configurations and under a variety of circumstances. Mock-ups of targets have been made in order to test computer vision systems, and autopilot systems are to be tested under varying levels of autonomy ranging from waypoint navigation to target location.

Before each flight, a preflight check will be conducted to verify that all systems are functioning normally. A flight plan document will be referred to for expected operations and remedies for wing control surfaces, power systems, and other electronics both on the plane and on the GCS. A pilot will be accompanied by a spotter with binoculars to verify the integrity of the airframe during flights. Other team members will be focused on operating the GCS and speaking out changes in status of component systems during flights.

For any unexpected deviation from the mission plan, the team captain has the final say in whether manual takeover is warranted. The most senior member of the aero-mechanical and electrical team will provide guidance on whether a given malfunction requires emergency landing.

4. Safety, Risks, and Mitigation

4.1 Developmental Risks and Mitigations

Each of the main developmental risks identified in this document had the potential to prevent the team’s competition in the AUVSI SUAS competition. The below table provides a brief summary of these risks.

Table 5. Potential Developmental Risks

Developmental Risk	Severity	Occurrence	Impact	Mitigation Strategy
Vehicle damage ranging from minor to total loss.	High	Rare	Loss of developmental time and money.	A reserve fund was established to replace any broken parts, or if necessary, the entire UAS. Additionally, strict flight protocols were enacted.
Inability to develop computer vision system by deadline.	High	Infrequent	Loss of ability to perform target recognition tasks.	Create timeline for computer vision system development, updating regularly throughout season.
Inability to complete a custom flight controller and autopilot system.	High	Infrequent	Degraded ability to meet competition requirements, and loss of time.	An alternative system based on the Pixhawk platform was developed in tandem with the custom system, in case of failure.
Inability to create a custom ground control station.	Moderate	Infrequent	Degraded ability to meet competition requirements.	A multitude of MAVLink-compatible ground stations are available that can be quickly integrated with other systems.

4.2 Mission Risks and Mitigations

Due to the one-time nature of the competition, malfunctions or mishaps may endanger the performance of the team. In addition, a failure to mitigate potential risks could endanger flight readiness review approval. Thus, a number of potential risks have been assessed and mitigated.

4.2.1 Loss of Telemetry, Manual and Autonomous Control

A loss of autonomous control would entail the flight controller restarting mid-flight or communications ceasing during flight. To mitigate this, flight controller models are stability tested before they are taken into the air. Changes to control are tested on the ground by tilting the plane by hand before they are flown.

In case of intermittent loss of telemetry due to power fluctuations or damage, the transceiver antennas are designed so that no matter the orientation of the plane the transmission pattern will remain strong.

4.2.2 Malfunctioning Control Surface/Motor

In case of control surface or motor not responding or responding inappropriately to input, all control surfaces are checked prior to flight.

4.3 Operational Risks and Mitigations

Due to the energy required to fly, there are numerous scenarios which could cause danger to crew or property. While preparing for every eventuality is difficult, there are some common scenarios that could cause a risk to Purdue Aerial Robotics’ property or members.

4.3.1 Uncontrolled Flight

One of the greatest operational risks is the possibility of uncontrolled flight. While this is a very rare possibility, it is an extremely severe condition. This is mitigated by the inclusion of a manual piloting system, which can be engaged at any time by the designated safety pilot.

4.3.2 Loss of Communications

A possibility that can occur during flight is the loss of communications between the safety pilot and the plane. To mitigate this, transmitters with a range of more than 1km have been used to ensure that communications will be stable at competition ranges. Communications are tested during setup to ensure that the UAS is sending and receiving information.