Design of the Minion Research Platform for the 2016 Maritime RobotX Challenge

Christopher M. Ash, Nate D. Bloom, Jeremy D. Brown, Abby N. Butka, Stephen P. Cronin, Grady C. Delp, Robert T. Goring, Christopher J. Hockley, Zachary R. Joswick, Diego Lodato, Nicholas R. Middlebrooks, Brandon F. Mathews, Joshua L. Pletz, Jefferson S. Romney, Marco A. Schoener, Nicholas C. Schultz, Alena M. Thompson, David J. Thompson, Johazais K. Wyble, & Timothy A. Zuercher
Dr. Charles F. Reinholtz, Dr. Eric J. Coyle, Dr. Brian K. Butka, Dr. Patrick N. Currier

***Abstract*— Embry-Riddle Aeronautical University (ERAU) has made significant improvements to their fully autonomous research platform, Minion. The new autonomy package, enables modularity in Minions mission, and can be configured to perform advanced perception, navigation, and decision making tasks. To complete these tasks, Minion uses sophisticated sensory and perception algorithms fusing the data from a suite consisting of four LIDARs, two wide-angle cameras, and a high precision GPS/INS. This data is fed into path planning and decision making algorithms which determine Minions actions. These algorithms include complex neural network visual detection and tracking, 3D Multi-Variate Gaussian classification, and dynamically updating path planning.**

Taking lessons learned from the 2014 competition and operations since, the Minion platform was developed emphasizing modularity, allowing it to meet the objectives of the 2016 RobotX Challenge and the demands of future research and teaching interests. This includes plug-in expandability through a unified communications backbone for all peripheral systems and custom designed wide-input power distribution and propulsion control systems.

These lessons also influenced the hardware design. The mounting system is based around the large underslung MAST system, and the modular payload tray which uses Picatinny rails and threaded attachment points. All of Minion's systems are rated to survive operations in adverse weather conditions, including high heat, high humidity, and heavy precipitation, and have been tested in these environments.

In the course of development, sub groups were formed and solutions to individual tasks were validated using simulations, recorded data, and over 100 hours of in water testing. The end result of this is an advanced platform, that is robust, reliable, and readily upgradable.

I. Introduction

A. Team and Vehicle Overview

Embry-Riddle Aeronautical University's (ERAU) Team Minion consists of students ranging from the 11th grade to Ph.D. candidates, with backgrounds in Software, Electrical, and Mechanical Engineering. The team has combined experience working on multiple autonomous platforms, including entries into the AUVSI Foundation RoboBoat and RoboSub competitions. Team Minion has also partnered with multiple organizations in the unmanned and autonomous systems industry. This experience has led the development of the second-generation Minion ASV, shown with the team in Figure 1.

Minion is designed to be behaviorally robust, rugged and readily upgradable to meet mission requirements. To ensure robustness, all components are designed to meet harsh environmental conditions including high heat, humidity, and precipitation. Hardware modularity is ensured through a common communication standard and a payload tray that includes Picatinny rails for easy and secure mounting. Wherever possible, hardware critical to operation has been mounted below the WAM-V's payload tray. The payload tray has defined predrilled pattern for the simple standardized mounting of additional payloads, e.g. UAV launch and recovery. Minion uses an Ethernet communication and power backbone allowing for the simple addition of external sensors, behaviors, and functionality, such as sub vehicles and turrets.

Figure 1: ERAU's Minion Team with the second generation Minion ASV

B. Software Overview

Software onboard Minion is broken into individual process modules. These modules execute in parallel and communicate asynchronously; using a publisher, subscriber messaging system. This enables modules to run at different rates and be selectively activated and deactivated improving overall system efficiency. The competition architecture is shown in Figure 2.

The Mission Tracker aggregates data from various modules and determines the best current objective to complete the mission. It communicates the objective to the Path Planner which calculates the optimal path to complete the new objective. The Mission Tracker also communicates the objective to the sensory modules, enabling and disabling processing algorithms based on the current objective. The function of each module is discussed in Vehicle Design.

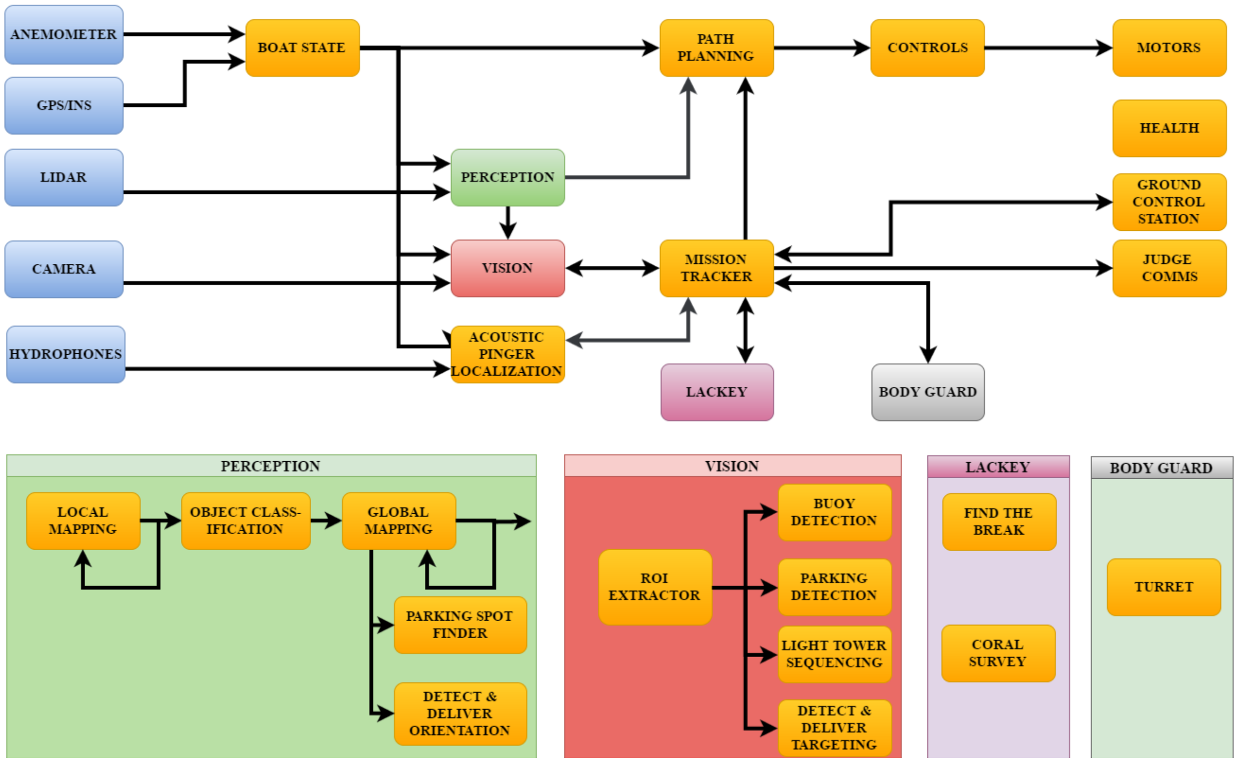

Figure 2: Minion ASV Software Architecture

C. Organization

Section II covers the teams design strategy and team structure. Section III discusses the hardware and software modules that make up the Minion platform. Section IV. presents simulation and testing results. Section V details the systems readiness and the team's future work.

II. Design Strategy

Team Minion made a fundamental shift in design strategy for the 2016 Maritime RobotX Challenge. After the 2014 competition the team pursued multiple student research projects and grant proposals involving the Minion ASV. Thus, Minion is designed as a reconfigurable sensor platform that is also autonomous, making it a more valuable tool for both research and teaching.

Team Minion is broken into six multidisciplinary groups. These groups are: Mechanical Systems, Electrical Systems, General Perception, Task Perception, Path Planning and Controls, and Ground Systems. The groups worked in parallel to create and improve each modular subsystem and coordinated as required with the project technical lead and other groups. Team leadership identified system requirements and strategies required to complete the competition tasks and fulfill the research objectives, which were given to each subgroup for implementation.

A. Mechanical Systems

The primary objective of the Mechanical Systems team was to ensure the reliability, modularity, and upgradability of the hardware system. The goal was to remodel the ASV infrastructure to include all basic hardware components under the deck to allow for mounting of sensors and modular mission packages.

In addition to the ASV infrastructure remodeling, a multi-purpose stabilized gimbal was created and outfitted with a pneumatic launcher (Bodyguard) to complete the Detect and Deliver task, and an unmanned underwater vehicle (Lackey) was designed to complete the Underwater Shape Identification and Find the Break tasks.

B. Electrical Systems

The Electrical systems group had three primary tasks. The first task was to enable Minion to operate at a higher voltage (50V nominal) than the 2014 vehicle could. This enables the vehicle to travel twice as fast and operate with double the battery storage.

The team was also tasked with replacing the onboard computing solution with a system that was more reliable, expandable, and powerful.

Finally, to increase reliability and to enable advanced control techniques the team was tasked with designing motor controllers for the Rim Driven Propellers (RDP).

C. General Perception

The task of the General Perception group was to design, calibrate, and implement software for object detection and classification using LIDAR and camera data. While the single LIDAR system used in 2014 was effective at close range, objects beyond 10m would often be missed due to gaps in the LIDAR coverage. The addition of new waterborne objects in 2016 lead the group to upgrade to four LIDAR sensors that were configured using simulation to enable the goals of reliable detection and classification of objects up to 25m from the vessel and representation of objects as geometric polygons.

In the 2014 RobotX Competition, the Minion ASV also relied on two webcams for vision sensing. These webcams used 20% of the available computing resources for image capture alone and led to significant errors in image construction, including over/under exposures, inconsistent distortion across the image, flaring from the built-in lens, and poor low-light images. Thus, the group sought to provide a clearer, less computationally intensive camera solution, providing higher fidelity images at a greater range and in a wider variety of lighting conditions.

The group was also tasked with classification of objects to isolate regions of interest for the Task Perception group, which reduces the image size and computation time for perception algorithms used in specific tasks. This approach also decouples the use of these different perception algorithms from the general perception system used for path planning and obstacle avoidance.

D. Task Perception

The goal of the Task Perception group is to identify features necessary to complete competition tasks, such as shape and color of docking symbols. Most tasks require different forms of vision processing. To decrease sensitivity due to changes in weather and lighting conditions, more robust techniques were explored including pattern matching, and deep learning. Pattern matching can enable strong shape recognition on small images. Using deep learning, varied lighting conditions can be handled by adding images of different lighting to the training set. Deep learning is more complex than traditional computer vision, however the process increases the reliability of the system for navigation around totems and detecting the light tower. For the camera onboard the submersible, a more traditional computer vision approach is used, since lighting can be more easily controlled using an onboard illuminator.

The group also needed to design and implement a hydrophone array to determine the location of an underwater pinger. In 2014, Minion featured a unique feathering hydrophone array, with two hydrophones mounted on each pontoon. This array design led to indeterminate solutions and unreliable pinger locations. For 2016 the team explored simpler ultra-short baseline arrays attached to a deployable arm.

To increase the accuracy of the acoustics, the team explored methods of reducing environmental noise such as turning the motors off during the hydrophone recording process. Signal processing techniques were also explored to reject noise and improve robustness.

E. Path Planning & Controls

The goal of the Path Planning & Controls group is to navigate successfully and accurately while accomplishing tasks that are reconfigurable in order and objective.

A new Mission Tracker architecture allows for missions to be dynamically reconfigured to add or reorder tasks while reusing base primitive behaviors to allow for easier validation.

With the increase in sensing horizon afforded to Minion through advances made by the General Perception group, a new mapping and controls paradigm was required. This is largely because the Dubbins approach used in the 2014 competition treated all objects as circles, while the new methods produce polygon objects. A more rigorous cell-based A* trajectory planning system was investigated as well as methods that account for vehicle dynamics.

The group also explored changing the control strategy to more accurately track paths than the 2014 version by implementing new algorithms enabling improved maneuvering in tight spaces such as docking or when near obstacles.

F. Ground Systems

To accurately simulate the course, the Ground Systems group was charged with developing surrogates for all course objects and features. This includes a dock, Detect and Deliver platform, Light Tower, Totems, acoustic pinger and obstacle field. The design of these elements had to be consistent with objects specified in the 2016 competition rules, but with the added goals of being inexpensive and feasible to set up and take down in daily testing. Thus, surrogates were constructed that provided the same sensor information to Minion, see Figure 3.

The group also needed to create a transportable ground station and judging interface. In 2014 transporting the ground station between fields was labor and time intensive. This was a key requirement for the 2016 system. Another key issue from 2014 was that the ground station was the only way judges could view data from Minion. The goals of the new system were ease of use, set up time, reliability, and judges access. Numerous solutions to this were investigated including a separate judges station and ground station, and unifying the ground station hardware. Finally, the group is responsible for vehicle and team logistics during testing, travel, and competition.

Figure 3: Simulated boat dock with Minion on approach.

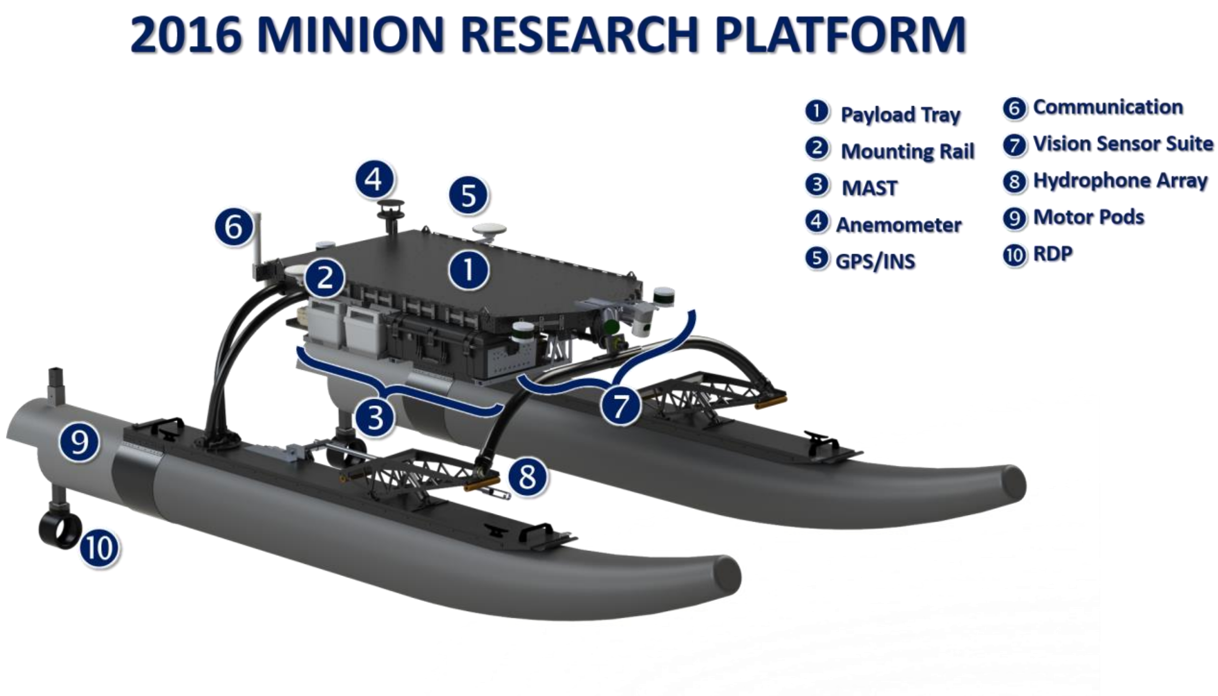

Figure 4. Minion ASV Design

III. Vehicle Design

A. Hardware Design

Hardware on Minion consists of commercial off-the-shelf sensing, as well as custom-built ruggedized computing hardware and power distribution systems. System propulsion consists of a pair of fixed electrically-powered RDPs attached to custom Motor Pods, see Figure 4. These thrusters minimize the potential for ingesting debris, maximizing safety to personnel and marine life. The RDP thrusters are driven using motor controllers designed and built in house which enable better control and monitoring of the motor systems. They also enable Minion to operate in both 25V and 50V nominal configurations. In order to eliminate potential points of failure incurred by a steering mechanism, Minion steers using the differential thrust between the two RDPs creating a yaw moment.

The MAST (Minion Autonomous Systems Tray) is the largest single structural change to the ERAU WAM-V platform. The MAST consists of a single support beam and mounting hardpoint for Minion's batteries, ESCs, control board, expansion bay, and computer hardware. These systems are readily accessible, upgradable, and removable using only a ratchet, and mount seamlessly onto the WAM-V with no modifications to the existing structure.

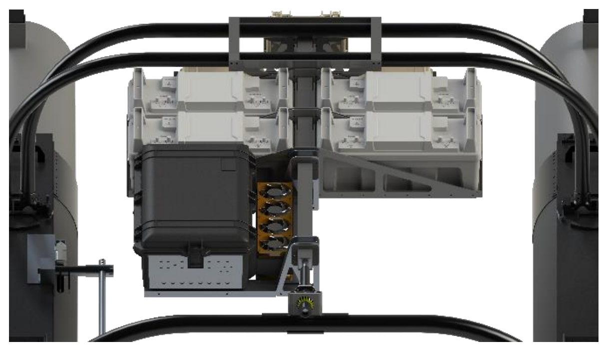

Figure 5: Minion Autonomous Systems Tray (MAST)

To achieve this, a single large beam was installed running the length of the payload tray, attaching to the rear arches and the forward revolute joint. This beam is sized to ensure safe carriage of Minion's hardware in both static and dynamic loading cases, which was then verified using dummy concrete payloads before installation of any hardware. The beam forms the backbone upon which battery and computing payloads are attached as well as the potential for additional mission specific configurations, such as the Lackey sub-vehicle. The MAST lowers Minion's overall CG by 12 cm and frees the payload tray from basic operation hardware such as the system computer and batteries to enable mounting of mission packages.

The Minion ASV's payload tray has also been upgraded with a modular rail system using the ubiquitous Picatinny mounting system to ensure easy, tool-less, and secure installation of hardware. Atop the payload tray is a pattern of rivetless nutplates with 20 cm spacing that covers a majority of the payload tray.

B. Subvehicle Design

Minion's subvehicle, Lackey (Figure 6), is designed to act as an autonomous mobile sensor platform. Based on the SWASH (Small Waterplane Area Single Hull) concept, it can operate with minimal interaction with the water surface, resulting in a more stable sensing platform. The 73.6 cm long by 12 cm in diameter hydrodynamically tailored hull is based upon the testing done by the U.S. Navy in the development of the USS Albacore and offers both minimal drag and high maneuverability without sacrificing internal usable volume [1].

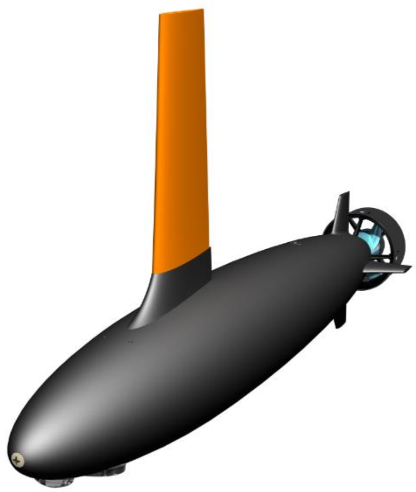

Figure 6: 3-D printed Lackey platform

In addition to stability, the mast provides a clear view of the sky for both GPS based sensing and normal RF based communications. Manufactured using the fused deposition modeling method of direct digital manufacturing, the nylon body is both tough and flexible and can be scaled and rebuilt as mission requirements dictate without the fabrication of new tooling equipment. In its standard configuration, Lackey utilizes an onboard downward facing color camera and high intensity illuminator, as well as a fully featured IMU/Magnetometer combination, single board antenna and Wi-Fi based data links.

C. Perception Software

1) LIDAR Object Detection and Classification

The Minion Perception suite uses four Velodyne LIDAR sensors providing 3.5 million data points per second. Prior to their selection the predicted LIDAR returns were simulated in various configurations. Ultimately, the configuration chosen result in the return pattern seen Figure 8. This pattern provides a high density of returns from competition objects. Returns from the water's surface are ignored using an intensity threshold, as water is an IR black body absorbing most of the laser light.

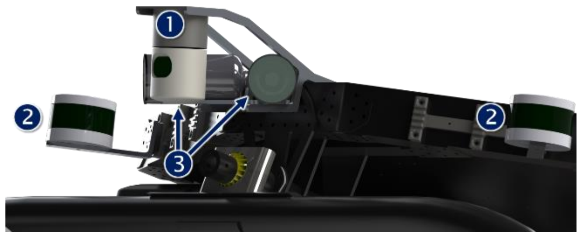

Figure 7: Bow Perception Suite with a center Velodyne HDL32-E (1), port and starboard VLP-16HD (2) and two Teledyne cameras (3).

The remaining LIDAR returns are mapped into the fixed NED reference frame the boat uses for navigation using the known mounting of each sensor and the measured state of the vessel. The set of NED points are quantized to a 10cm resolution and used to fill a 3D occupancy grid. This voxel grid is referenced to a fixed frame and these points are retained for a period of time (four seconds typically), without the need to re-compute point locations as the vessel moves. The voxel grid is then flattened vertically, to yield a 2D occupancy grid of filled cells on the same plane as the water's surface.

To extract objects, the perception system uses the novel approach of treating this 2D grid as a binary image from which image contours can be extracted, resulting in a list of polygon object bounds with NED frame vertices. Only polygons found in the region where high point density is expected, plotted as the black lines in Figure 8, are retained from this process and merged with previously known objects outside this region. This methodology was first developed and verified using the LIDAR simulation environment before porting to the Minion platform. This process has shown to be highly accurate at identifying objects in an operating area, as shown in Figure 9.

The voxel cells and polygon vertices of each object are used to identify object features of length, width, height, perimeter, 2D area and surface area, which are then passed to a Multi-Variate Gaussian (MVG) classifier. This classifier estimates the probability the observed object is each of the known competition elements. If the highest probability among the considered classes is above a tunable threshold, then the object is labeled as part of this class, otherwise it is labeled as unknown.

Once the object has been detected and identified based on its spatial properties, the list of objects and the object class is passed to the cameras for identification of visual properties, such as color, light sequencing and shape detection. To ensure proper decision making, object class is not considered final until 90% of the classification results for the object are from a single class over a 6 second period with the system running at a rate of 5Hz.

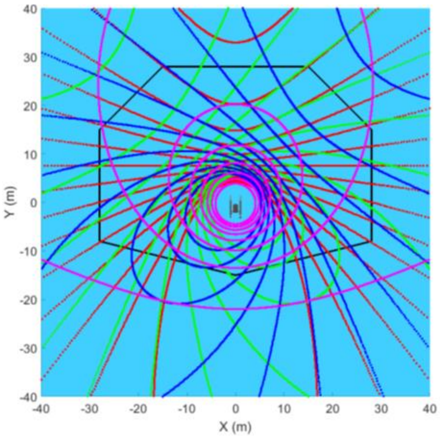

Figure 8: Predicted Lidar returns on the water surface. Black – Region with high return density, Blue – Port VLP-16HD returns, Red – Bow HDL32-E returns, Green – Starboard VLP-16HD returns, Magenta- Stern VLP-16 returns.

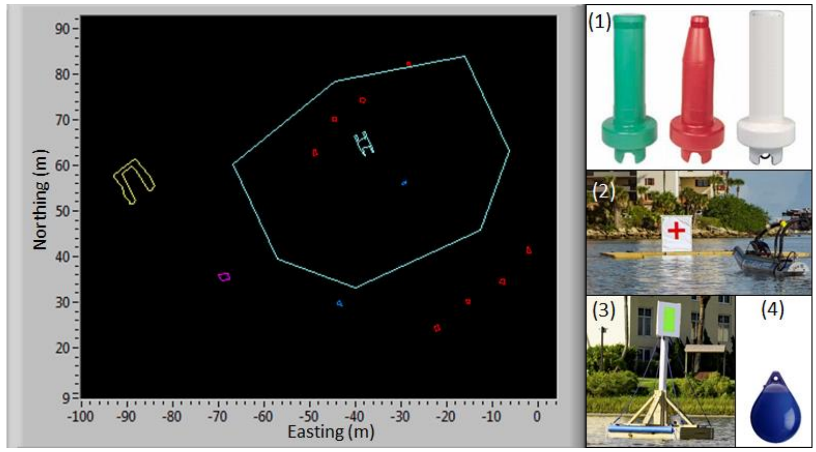

Figure 9: Real-time classification results. Red polygons are (1) Taylor-Made Buoys, yellow polygons are (2) the floating dock, magenta polygons are (3) the light tower, and blue polygons are (5) Polyform buoys. The cyan region forms the boundary used to map previously known and newly detected objects.

2) Camera Perception and Classification

The vision suite processes image streams from two Teledyne Dalsa Genie 5 megapixel cameras that minimally impact the CPU. The cameras have 100° field-of-view and 10° overlap between the images. To negate the impact of ambient light the cameras are equipped with both polarizing and IR filters.

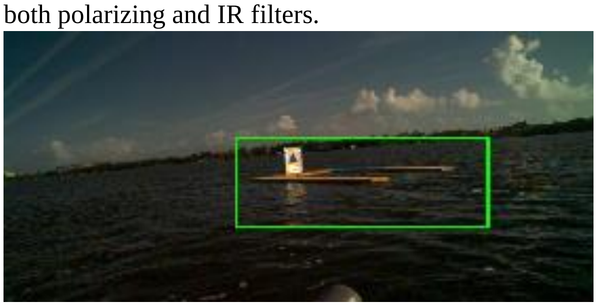

Figure 10: An example of a full image with the detected region of interest highlighted by a green box

Utilizing known camera parameters such as mounting position, orientation, focal length, and distortion, an object's NED position can be accurately converted into the pixel frame of the camera, enabling extraction of a region of interest (ROI) for an object detected via LIDAR. This improves the efficiency of computer vision systems by reducing image areas that may cause false positives. An example of an extracted region of interest is shown in Figure 10 where the region is highlighted as green box for visualization. From the list of objects previously discussed, regions of interest are passed to individual modules for feature identification.

3) Vision Application to Competition Tasks

a. Docking and Detect and Deliver Shape Detection

For the docking task Minion needs to identify the shape and color of the dock target. Using the appropriate ROI, the camera image is post-processed with a color threshold, removing colors outside of set HSL ranges, yielding a binary image as shown in Figure 11. The resulting binary image is then compared with a shape classifier to generate a score for that color and shape combination. This loop is iterated on the original color image for every shape and color combination possible with the dock signs, with the highest scoring output classified as the accurate result. The detect and Deliver challenge uses the same approach as docking task as the shapes and colors are identical and the target is found using the LIDAR perception software.

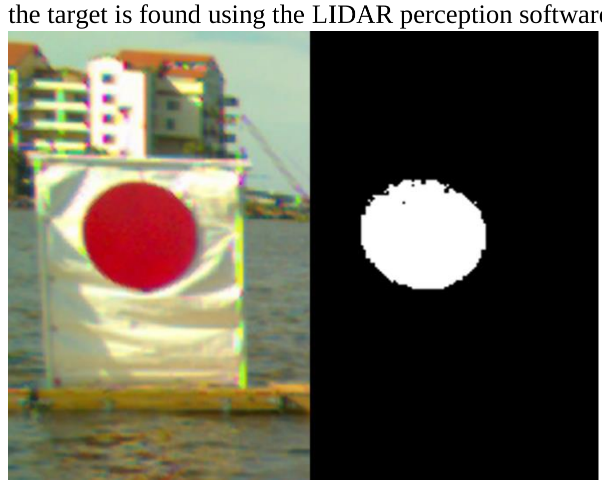

Figure 11: A dock sign image returned from the vision module (left) compared to the color threshold result (right)

b. Scan the Code Light Sequence

For the Scan the Code task, a sequence of colors must be read from a tower with a color illuminated LCD screen as shown in Figure 12. The cameras via their respective ROI captures the colors and process the images using a deep learning Faster R-CNN algorithm. Faster R-CNN was chosen due to its speed and accuracy when performing object identification and classification. Faster R-CNN can identify multiple objects per frame, as well as identify objects partially in frame. One of the hardest aspects of vision processing is consistently being able to classify colors in various lighting conditions. This method can classify the color of the light panel without needing to detect the tower itself. Experimental testing has shown the panel color can be correctly detected with various lighting conditions.

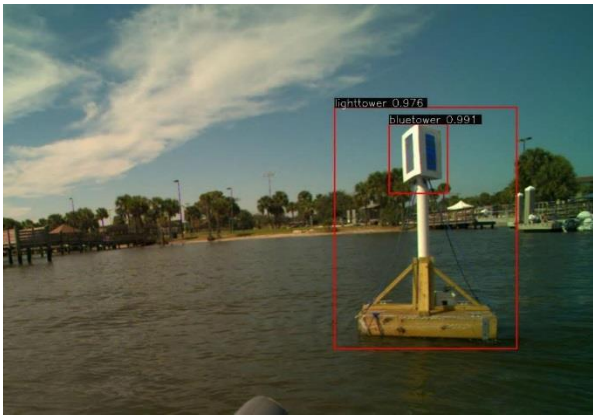

Figure 12 The Scan the Code Module detecting the color emitted from the light tower (blue) using Faster R-CNN algorithm with a 0.991 confidence.

The Faster R-CNN algorithm is based on the Caffe framework. This technique is feasible due to the system's two GTX1080 graphics cards which allow for large memory and processing requirements of the neural network.

c. Find the Totem

The totems are initially identified via the LIDAR perception module generating an ROI and then the color determined using the same Faster R-CCN algorithm as with the Scan the Code Light Sequence. However, a different network is trained to determine the individual colors of the totems.

d. Find the Break

To find a break in a series of underwater colored pipe lines as required in competition, images are extracted from Lackey's underwater camera. The general perception software indicates the area to scan by identifying the two sight line markers. A two pass approach is used to first confirm the direction of the underwater lines, and second scan for breaks.

The images are processed to extract the underwater markers by converting them to the HSV color space and applying a hue threshold tuned to the orange and yellow pipe line colors. Contouring is then used to extract the lane markers and their orientation. Confirming the orientation serves to further increase confidence that the correct color marker has been identified. From this point, objects are tracked between frames to build a complete understanding of all markers. With the position of the markers determined, a simple algorithm is used to determine where orange markers are between yellow markers, indicating a break.

D. Acoustic Sensor Processing

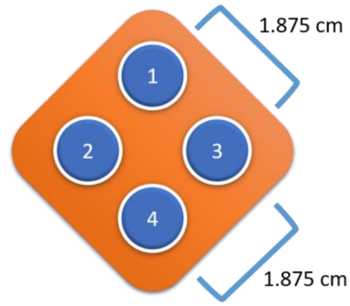

Figure 13 Hydrophone Array with four Teledyne Reson Sensors.

The primary component used in the hydrophone task are four Teledyne Reson sensors arranged in a square array, with each sensor placed less than half the wavelength of the highest frequency ping, 40 kHz. The signal measured by the hydrophones are then amplified and read using a NI-9222 DAQ. This amplification is accomplished using a custom circuit board, which also contains an embedded processor and supporting circuitry to raise and lower the arm on which the hydrophone array is mounted.

The hydrophones are a planar array, which allows for a heading to the ping source to be determined. The algorithm used for processing the hydrophone data is a multi-step process for conditioning and filtering the data. It begins with a band-pass filter around the target frequency. This was determined to be more computationally efficient over a Fast Fourier Transform. After filtering, the front of the waveform is detected using level detection. This method isolates the front part of the waveform which has not come through a reflection off another surface, which results in phase shifts of the data.

With the start of the waveform detected, 200 samples are taken for processing. This number of samples is ten full phases at the slowest frequency in consideration, 25 kHz. The major factor considered in signal corruption is the high sensitivity of the hydrophones to motor noise and the proximity of the frequency of the switching motor controllers to the pinger frequencies. As such, the Mission Tracker instructs Minion to drift during the sampling of hydrophone data. To process the filtered data, a pure geometric solution is employed based the time difference of arrival to the sensors. The heading is converted to a global coordinate frame and output to the Mission Tracker.

E. Path Planning and Control

1) Mission Tracker

The Mission Tracker module handles the high-level decision making necessary to accomplish the required tasks. This module uses a subsumption architecture to activate individual task behaviors as defined by an XML format file that defines overall mission goals and task sequencing. An overall progress tracker monitors mission state and launches or terminates behaviors to ensure continued mission progress by either accomplishing goals or timing out.

Individual task behaviors are implemented using a common format that takes inputs from the State, Vision, and Hydrophone modules to determine the state of the necessary elements for the task. The task then implements primitives such as "drive to waypoint," "locate gate," "search for target," or "circle target" to accomplish the task. This approach allows for efficient code reuse and validation of individual components, both of which were problems in the 2014 RobotX Challenge.

2) Path Planner

The Path Planner module handles basic navigation and obstacle avoidance for Minion. It receives four degree of freedom (northing, easting, heading, speed) target waypoints from the Mission Tracker and input from the Perception module in the form of obstacle locations and boundaries.

The Path Planner operates a cell-based 4D (Northing, Easting, Heading, Velocity) A* search to find an optimal path through the known obstacle field from its current pose to the pose required by the 3D target point. To enable Minion to dynamically grow its search grid the search is rotated and translated into a goal oriented frame so that the target point always occupies (0,0,0,0). This enables significant computational savings in the heuristics and path cost calculations. The nodes are dynamically generated as the path is explored enabling memory savings over a statically declared grid. When a unique northing easting pair is explored the grid cell is checked against the obstacle locations and if they are within a user specified distance of the grid cell center the cell is considered occupied and not traversable. This enables growing the obstacles to ensure sufficient clearance for the width of the WAM-V.

3) Vector Control Module

Paths are sent from the Path Planner to the Vector Control Module as a series of connected line segments. The VCM's task is to follow the path as accurately as possible until a new path is received. Two different algorithms were investigated to enable path following.

Both algorithms generate target speeds and yaw rates, which are tracked using PID controllers. Each PID incorporates a feed-forward term to increase accuracy and has limits to prevent integral windup.

The first algorithm rotates the boat into a coordinate system aligned with the path. It then calculates an error along the path e_{FB} and an error orthogonal to the path segment e_{CT}. Forward velocity, v, in the boat frame is then controlled to keep the vessel at the required velocity to achieve the trajectory:

$$v = v_{DES} + k_1 e_{FB}. \quad (1)$$

The desired yaw rate, ω, is then chosen to attempt to achieve the desired heading of the path segment while minimizing the cross-track error:

$$\omega = k_2\left(\theta_P - \theta_B - \text{atan2}(v_Y, v_x) + k_3 \tanh(k_4 e_{CT})\right) \quad (2)$$

Note the use of the heading of the actual velocity vector to allow the vessel to better account for disturbances such as wind or current by "crabbing" in the disturbance at an angle. The cross-track error and path heading are calculated using the point on the path closest to the present location of the vessel.

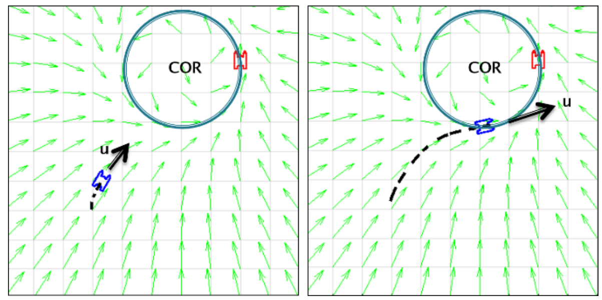

Figure 14 Boat (blue) driving to the desired Boat (red) using velocity vector field technique with Center of Rotations (COR).

A different control strategy uses a velocity vector field to consistently account for external disturbances to the Minion ASV [2]. The velocity vector field is derived from using the desired position (X_r, Y_r) heading (θ_r), velocity v_r and angular velocity ω_r to calculate a center of rotation. The ASV 's desired heading angle, θ_d, is toward the circle when distance is large and converges tangent to the circle when distance is small as shown in Figure 14.

$$v = \frac{Dv_r}{r}\cos(k_v e_\theta). \quad (3)$$

$$\omega = -k_\theta e_\theta = -k_\theta(\theta - \theta_d). \quad (4)$$

For this controller to efficiently drive the boat, the path is assumed to be generated in front of the boat and that no extreme corner turns are given in the course.

In-water testing has shown that the first algorithm produced better performance, resulting in its selection for competition.

4) Path Planner and Control Application to Competition Tasks

Driving around totems will be conducted by selecting target points in the Mission Tracker. The Mission Tracker will calculate a circle around the totems and then send four progressive target points to the Path Planner, causing it to execute a circle while still allowing the Path Planner to compensate for and avoid obstacles.

Due to the under actuated physical motor layout, true station keeping (x, y, heading) is not possible on Minion. Tasks requiring station keeping will be conducted using a heading hold mode in the Path Planner. The planner will request a zero speed and set a target heading. The speed control and yaw rate controllers in the control module will act to hold the speed and heading. The Mission Tracker will monitor the position of the vessel. In the event of excessive drift (>10m), the Mission Tracker will change the Path Planner back to navigation mode and return to the stationary waypoint.

F. Monitoring and Reporting

Team Minion's ground station is designed to monitor the vehicle's status and environment and sends the vehicle its tasks for each competition run. The goal of the ground station design is to quickly and reliably check the status and run-time settings of the vehicle in-real time with a professional, streamlined interface. This communication, along with all internal communication on the vessel and communication to the reporting tablet (discussed below) are streamlined using the Minion Core messaging protocol. The Minion Core messaging protocol is a publisher subscriber system developed by the team to easily facilitate the development and integration of new software modules into the communications network.

The ground station software, which is for team use separate from the task reporting tablet, is segmented into three sections: the map interface, the mission tracker, and the boat system monitor. The map interface uses a cached offline map image and displays the vehicle, course elements, and the vehicle's desired path in latitude and longitude space. The objective tracker takes the settings for the current competition run and communicates that over to the vehicle to perform when the run starts. The system monitor is responsible for overseeing the status and data from the vehicle's sensors and algorithms.

Figure 15 Minion Planner map interface to visually track its location

The Judges Tablet is a more transportable monitoring system for the judge's use at competition. The tablet displays all task reporting for the judges to view the status of the mission objectives along with features for investigating vehicle and sensor status (with GUIs for info and data from sensors), mission run time and vehicle uptime, and a touch screen map interface to follow the vehicle on the cached map system. More details in Appendix A.

IV. Experimental Results

Team Minion used a three-tiered validation system for the platform software and sensors. The team started by developing approaches using simulation, then moved to recording data and ensuring algorithm effectiveness on recorded data, and then moved to full system testing.

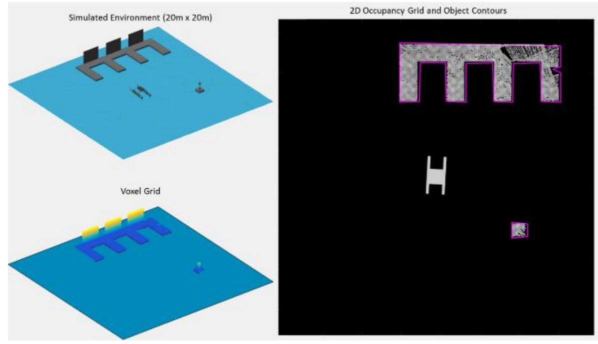

Figure 16 Simulation of Lidar returns on the Minion ASV. The voxel grid created from the Lidar returns is shown in the lower left, and the corresponding 2D occupancy grid and object contours are shown on the right.

A. Software Simulation

To tune and develop the LIDAR perception system, the team developed a MATLAB simulation environment that could give simulated LIDAR returns from any three-dimensional object created using CAD software. The simulation environment, shown in Figure 16, modeled Minion's dynamics as well as the error expected of the LIDAR sensors, and the GPS/INS system.

The dynamic model of Minion was also used to verify the controls approaches previously discussed. This boat model accounts for the hydrodynamics of the boat, which shows the drift and reaction forces when under power from the propulsion motors. The controls model was also used to determine the maximum complexity of the paths it is possible to follow and to troubleshoot errors in the control system in conjunction with experimental testing.

B. Data Playback

As part of the Minion Core protocol, the raw value of all sensor data is recorded. This enables playback of these messages to the boat modules for offline tuning and verification. This was particularly important for the vision and hydrophone team due to not having a proper method of simulating all the environmental factors that would affect these algorithms and approaches. Using recording in this way also enabled full offline debugging of system software including the internal passing of data between modules.

With approximately 50 hours of camera footage the cameras were experimentally calibrated to the current polarization and IR filter setup. The results from the tuned cameras yielded large amounts of data for testing the Scan the Code, Totem Recognition and Detect and Deliver. This yielded over 90% accuracy in reading the light tower, totem recognition and the detect and deliver obstacle. The dock bay additionally can reliably be detected with orientation and heading to less than 5 degrees and perform tracking on the docking symbol once identified.

After 22 hours of hydrophone testing, the algorithm was proven accurate to within ±2 degrees of the correct angle to the competition pinger. Testing in this manner ultimately showed the need for adequate noise filtering and the need to shut off the platform thrusters as previously stated.

C. In-Water Testing

The Embry-Riddle Minion Team had a rigorous testing schedule of deploying the Minion Research Platform in Central Florida's Halifax River at least once a week after hardware integration was complete. The boat performed over 100 test hours from 39ºF to 107ºF, pouring rain, and 20 knot winds.

During these tests, the vehicle's hardware was stress tested while data was recorded from the custom built course objects, including dock, light tower, and detect and deliver. Field testing was also used to tune algorithm parameters and ensure these procedures were well-defined.

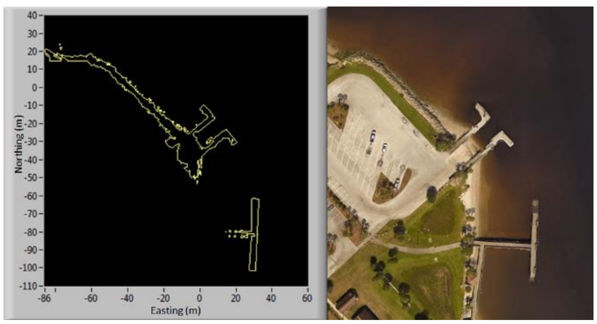

Figure 17 Real-time detection of pier, dock and shoreline contours compared to satellite imagery.

The LIDAR classification module was found to be 93.5% accurate in testing classifying the five spatially distinct competition objects (Taylor-Made Buoys, Polyform Buoys, Detect and Deliver Target, Docks, and the Light Tower), with example results shown in Figure 17. This accuracy has since been further increased by incorporating the previously stated time-based filter to finalize object class.

The team also completed and tested the 3-D printed submarine underwater and movement of the submarine via remote control.

V. Conclusion and Future Work

ERAU has developed the Minion ASV research platform with the twin goals of competing in the Maritime RobotX Challenge in Hawaii in 2016, and serving as a research and teaching tool for the future. The Minion research platform incorporates a novel hardware and software design, heavily leveraging the ideas of interchangeable hardware modules. Adaptable sensing modalities and hardware are key to this approach, coupled with advanced perception and Guidance Navigation and Control algorithms. A series of ERAU platform firsts were achieved during this development process, including the development of an advanced neural network image capture tracking and classification software suite and a voxel based Multi-Variate Gaussian classification and mapping based software package. This classification paradigm forms the cornerstone of the perception suite which has proven over 90% accurate at ranges extending out beyond 25 meters. This perception suite as a whole enables Minion to complete a variety of maritime autonomous navigation tasks. In order to better meet the modular design ethos, custom design power distribution systems were developed for the platform, deviating from the single board used in the 2014 platform. This ensured that failure on one line or component did not jeopardize operation of the system as a whole.

Leveraging the understanding of maritime robotic systems development and operations gained from previous years RoboBoat and RobotX an entire new group of students have been trained. The scope of the project was also larger than any previous effort and required a much higher degree of technical competence and logistical work than anything ERAU has previously done. With this system the team believes that it can successfully complete the tasks in the 2016 Maritime RobotX Challenge.

Future plans for Minion involve using it as a part of a multi-agent autonomous system, which was started with the integration of the Lackey platform. Deploying smaller unmanned aircraft such as the VTOL Flying wing platform or submersibles to widen its sensor horizon. Additional sensor modalities are also being researched for additions to Minion, including Linescan sonar systems and radar. Missions for such a system include environmental monitoring, wildlife protection and harbor security.

VI. Acknowledgements

The ERAU Minion Team would like to acknowledge the RobotX sponsors, AUVSI and ONR. The team also wishes to acknowledge its sponsors: Glenair, Velodyne Lidar, TORC Robotics APlus Mobile, Solidworks, Mathworks, ERAU COE, Motorolla, Rockwell Collins, ERAU SGA, Carlisle Interconnect Technologies, Stratasys, Advanced Circuits, Novatel, Pelican, Mean Well, IDEC, APlus Mobile, Boeing, Tru-Vu, AAI, Sparton NavEx, Brunswick Commercial &Government Products, and Web Daytona. Finally the team would like to thank all those individuals who have helped us in this project, without your assistance this project would not have been successful.

VII. REFERENCES

[1] The U.S. Navy, "Some Aspects of Submarine Design Part 2: Shape of a submarine 2026," The U.S. Navy, [Online]. Available: http://www.sname.org/HigherLogic/System/DownloadDocumentFile.ashx?DocumentFileKey=d00fc941-9b95-44b9-a5bc-854f4d6518eb.

[2] C. J. K. a. D. C. J. W. Kwon, "Vector field trajectory tracking control for wheeled mobile robots," in *Industrial Technology, 2009. ICIT 2009. IEEE International Conference on*, 2009, pp. 1-6.

[3] M. Gertler and L. Landweber, "Mathematical formulation of bodies of revolution," DTMB Report 719, Bethesda, 1950.

[4] M. Gertler, "Resistance experiments on a systematic series of streamlined bodies of revolution for application to the design of high speed submarines," DTMB Report C-297, Bethesda, 1950.

[5] "Model S," Tesla Motors, 2016. [Online]. Available: http://www.tesla.com/models.

[6] G. Papadopoulos, H. Kurniawati, A. S. B. M. Shariff and L. J. Wong, 3D-surface reconstruction for partially submerged marine structures using an autonomous surface vehicle, IEEE Intelligent robots and Systems, 2011, pp. 3551-3557.

[7] C. L. Kennedy, "Development of an Exteroceptive Sensor Suite on Unmanned Surface Vessels for Real-Time Classification of Navigational Markers," Daytona Beach, FL, 2014.

[8] "2016 Maritime RobotX Challenge Preliminary Task Descriptions," 3 November 2016. [Online].

[9] B. J. Guerreiro, C. Silvestre, R. Cunha and A. Pascoal, "Trajectory Tracking Nonlinear Model Predictive Control for Autonomous Surface Craft," *IEEE TRANSACTIONS ON CONTROL SYSTEMS TECHNOLOGY*, vol. 22, no. 6, pp. 2160-2175, November 2014.

[10] S. Ren, K. He, R. Girshick and J. Sun, "Advances in Neural Information Processing Systems," in *Faster {R-CNN}: Towards Real-Time Object Detection*, 2015.

VIII. Appendix

A. Situational Awareness

Autonomous systems require an enormous amount of information to operate, including road data, traffic and weather conditions, and information of their surroundings. The majority of this data is acquired through sensor suites, typically consisting of cameras, RADAR, and LIDAR. Up until recently, the hefty asking price for these sensors has limited their viability in a consumer market, relegating them to existence in government and university research labs, such as those competing in the Maritime RobotX Challenge.

The state of affairs presented above is no longer the status quo; driving-assistance systems are making their way into greater and greater numbers of consumer automobiles. Models from dozens of manufacturers offer upgrade packages for adaptive cruise control or lane-keeping assistance. Even this won't be the status quo for much longer; sensors, like all technology, are becoming faster, more advanced, and more affordable every year. In some vehicles, the necessary sensor suites required to support full autonomy can be purchased for less than $10,000 USD. This cost may price some consumers out of the autonomy market for some time, but it is not unreasonable for upgrades to high-end luxury cars to be in this range. Full-fledged autonomy in consumer vehicles will certainly be an early-adopters market for some time, but it may only be a few years before such sensors could become standard in certain high-end models.

With autonomous vehicles beginning to make it into the hands of consumers and commuters, a discussion must be held on how they behave, and how they communicate their behavior to their passengers. The presence of autonomous vehicles on the market will mean nothing if their systems are not trusted well enough to be adopted. Team Minion recognizes this challenge, and has found a novel approach to solving it.

To provide ourselves with an idea of what our vehicle is doing, Team Minion utilizes two separate ground-station interfaces. One of these interfaces, MinionTab, is tailor-made to fill the shoes of the Judges' Interface, as prescribed in the Task Descriptions document. In addition, this device provides an overview of the system and subsystems on the Minion ASV. Data can be tracked in real-time from any of our sensors, and the vessel itself can be localized on a map of the surrounding area. Two key indicators in the MinionTab interface that provide users with ASV situational awareness are the Mode and Current Task displays. The mode display indicates whether the vessel is operating in RC or Autonomous mode, or whether it is Emergency-Stopped. The Current Task display indicates which of the 8 RobotX Challenge Tasks the Minion ASV is currently engaged in.

This hand-held, tablet-based interface is complemented by our second, and more feature-heavy, ground station interface known as MinionPlanner. MinionPlanner provides a more detailed overview of what the boat is seeing and doing. In addition to featuring the same sensor data readouts and

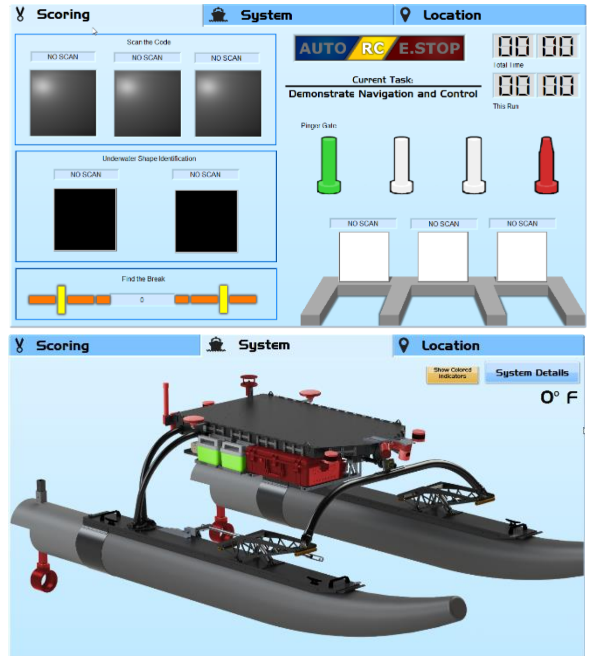

Figure 18: System Overview on MinionTab

localization features present on MinionTab, it can provide a visual of the objects Minion sees, and how it classifies them. It also displays the vessel's planned path, around obstacles and between challenges, and how closely the vessel is adhering to said path. Lastly, MinionPlanner provides actual input to the vessel that the user may utilize to affect the vessel's behavior.

These two interfaces could be seen performing dichotomous roles in a sort of user/developer mode relationship. Much like Microsoft provides Insider Preview builds to its many consumer software products and Google provides a Developer mode to enable certain debug features in Android, so too should autonomous consumer vehicles provide multiple interfaces for the user to interact with. In most cases, a user will only require the ability to see the operational mode, input a destination, and interact with the typical vehicle amenities, such as A/C and radio. If desired, however, it would not detract from the user experience to be provided with information such as the vehicle's planned route and how its sensors interpret other vehicles and objects nearby. A single instance of a user being able to see how accurately the system detects its' surroundings could go a long way towards earning that user's trust.

We find that this system-of-systems approach works best for autonomous vehicle user-interfaces. It is important not to provide users with so much information that they are drowned by it, yet allow them to view it if they wish to delve deeper. Transparency is the best way to earning user trust, and providing a dichotomous user-interface is the best way to provide transparency.