



*Cornell University Autonomous
Underwater Vehicle Team*

Fall 2016

Training Competition

Technical Report

Noel Picinich (nmp53)

December 10, 2016

Contents

1	Introduction	2
2	The Competition	2
2.1	Task Sequence	3
2.2	Scoring	3
3	Implementation	4
3.1	Pipes Vision	4
3.2	Bins Mission	5
4	Testing	7
4.1	Pipes Vision	7
4.2	Bins Mission	7
5	Results	7
5.1	Pipes Vision	7
5.2	Bins Mission	8
6	Conclusion	8

1 Introduction

At the start of the Fall 2016 semester, the six new software members were split into two teams and tasked with implementing a series of operations that Thor, the team's main sub, will perform in Teagle Pool's underwater arena. The goal of the annual training competition is to equip new members with an understanding of the vision and mission systems, as well as the ability to navigate and employ CUAUV's expansive software stack. Both teams, Kathode Ray Tube and The A Team, were challenged with writing vision and mission code to successfully perform a sequence of tasks detailed below. The following report describes components of The A Team's implementations including the pipe vision module, the bins vision module, and the bins mission. For more details regarding the other elements of The A Team's project, please refer to technical documentation by Evan Zhao and Horace He.

2 The Competition




Figure 1: Training Competition Course

2.1 Task Sequence

The following description corresponds visually to Figure 1. At the start of the competition, Thor is positioned at the East end of the pool with an orange and yellow buoy in view of the forward camera. The first task is to ram the orange buoy, back up, and pass over or around the buoy to complete the next task. The presence of the yellow buoy is to enforce color recognition in the buoy vision module. After ramming the buoy, Thor must align with the first orange pipe and travel in the direction the pipe points. If successful, the second orange pipe will come into view of Thor's downward camera and cue the subsequent task. The next task is to repeat the pipe alignment and travel towards the bins task. The bins challenge entails recognizing two bins on the arena floor, one indicated with a lightning bolt shape and another with a soda can shape. On the day of the competition, the judges tell the competitors which bin is the target and with a small alteration of the code, Thor should be able to drop markers into the specified bin.

2.2 Scoring

Each task completed successfully will be awarded a specified number of points. Tasks only partially completed or completed incorrectly may be awarded points at the discretion of the judges. There are a possible 500 static judging points which may be awarded based on competitor responses to questions about the design and implementation of the competition software. A specific point breakdown can be seen in Table 1.

Table 1: Scoring Breakdown by Task

Task	Points
Buoy Ram	500
Pipe: Align Position	250
Pipe: Align Heading	250
Bins: Drop On Target	1000
Bins: Drop On Incorrect Target	500
Static Judging	Up to 500

3 Implementation

In order to successfully complete the sequence of tasks, the A Team programmed a vision module and mission for each subsequent challenge. Vision modules were created by employing OpenCV, a library of computer vision functions, as well as CUAUV's previously implemented functions for displaying filtered vision on a Graphical User Interface (GUI). The mission code also employs CUAUV's previously written code for moving the sub, acquiring values from the sub's video input, targeting specific points on the pool floor, and several other helper functions. The values obtained from the vision module are communicated to the mission code by updating variable values in the team's shared memory system (SHM). This documentation will describe the specific details of how the pipe vision, bins vision, and bins mission were created. The A Team's full compilation of code also includes buoy vision code, buoy mission code, and pipe mission code that are included in the technical documentation of Evan Zhao and Horace He.

3.1 Pipes Vision

After completing the buoy bump, the pipe mission is coded to move straight until the orange pipe is in view of the sub's downward camera. As the sub drives forward it actively seeks to recognize rectangular shapes in view. In order to avoid recognizing the rectangular pool lane lines, filters were placed on the inputted vision that filter out items in view that aren't of the desired color. The LAB color space, which considers the luminosity and the a and b color channels, was chosen over other color spaces due to its color enhancing abilities when manipulating diluted or bright images. This advantage is an important ability for the sub's vision to have since the annual TRANSDEC competition CUAUV competes in is held in an outdoor pool where strong sunlight can greatly affect the sub's view. By manipulating the LAB components in the web GUI, ideal values were determined for obtaining the most clear image. Since the orange color of the pipe corresponds to certain LAB values, specific ranges were created that account for dilution of color but ultimately describe values representative of the orange pipe color. These ranges were used to isolate the orange pipe in the vision and create a black and white image of the rectangular pipe. Noise in the image corresponds to values that fell inside the range but were not part of the desired pipe shape. This noise was eliminated by employing the OpenCV erode function.

Since the sub may see the orange buoy and the orange pipe, the pipe

vision had to be able to differentiate circular orange objects from rectangular orange objects. OpenCV functions were used to draw contours along the black and white boundary, and then determine the minimum area rectangle enclosing the drawn contours, as well as the area the contours enclose (refer to Figure 2¹). The more rectangular a shape is, the closer the contour area will be to the area of the minimum area rectangle. The ratio of these two values was assigned to each contour drawn in the vision. The list of each contour was further filtered by removing contours with an area less than 200 and a ratio less than .6. The contours meeting these requirements were then drawn in the vision using OpenCV's draw contours function.

The final step in implementing the pipe vision was changing values in SHM so that the pipe mission could have access to the center coordinates of the perceived pipe as well as the angle of the pipe to the sub's current heading. The A Team used CUAUV's "get angle from rotated rect" function and assigned the returned value to the pipe results angle variable in SHM. The x and y coordinates of the center of the rectangle were acquired through the OpenCV minimum area rectangle function and assigned to the pipe results center x and center y variables.

3.2 Bins Mission

The goal of the bins mission is to go forward until a Bin is found, center the submarine at the true center of the bin, align the submarine such that it is parallel with the shorter edge of the bins, go to a greater depth, and drop markers onto the bin.

The first step in the bins mission is setting the x velocity of the sub to 1 so that it drives forward, while the bins vision is actively working to identify bin objects on the pool floor. When a bin is found, the bins vision will set the SHM yellow bin probability variable to one which the bins mission code continuously checks. When the probability is set to 1, the bins mission knows that the sub has found the bins and stops moving forward by zeroing the velocity. The next task in the bins mission sequence is to center over the desired bin. This step relies on the bins vision to set SHM yellow bin center x and center y variables to the coordinates of the center of the correct bin. The A Team's bin mission passes these values into a CUAUV "downward target" function which takes in a specific coordinate and the coordinates of the center of the sub's downward camera view to translate the sub's position to the desired point.

¹OpenCV Documentation - Contour Features: http://docs.opencv.org/3.1.0/dd/d49/tutorial_py_contour_features.html




Figure 2: Determining Rectangularity: The red rectangle represents the minimum area rectangle around the lightning bolt shape. In determining the rectangularity of the lightning bolt if placed in view of the sub, the A Team’s pipe vision would draw contours around the edges of the lightning bolt, determine the area within these contours, draw the smallest possible rectangle enclosing the shape (shown in red), find the area of this rectangle, and finally, calculate the ratio of the area of the bolt to the area of the minimum rectangle. This ratio, a value between 0 and 1, would express the rectangularity of the observed shape where the closer the value is to 1, the more rectangular the object is.

Once the sub is centered over the desired bin, it proceeds to the align task. In this task, the A Team utilizes CUAUV’s “relative to initial heading” function and inputs the sub’s relative angle from a value set in SHM by the vision module. Unlike in the pipes mission, the direction in which the sub aligns to the bin is arbitrary because the sub no longer has to proceed to another task. After aligning with the bin, the sub must increase its depth in order to have a greater probability of dropping the markers inside the bin shape. It is relevant to note that the bins referred to are actually flat boards on the pool floor with an inner rectangle representing the bottom of a bin and a white border that can receive a smaller fraction of points if contacted by the markers. While the sub lowers towards the bin, the A Team bin mission causes the sub to actively center over the bin. Once the sub is at the specified depth, the bin mission code call’s CUAUV’s “fire actuator” function by passing in the key phrase “both markers” that will cause the sub to drop two markers on the bin.

4 Testing

4.1 Pipes Vision

The pipes vision was tested in the vision GUI using previously recorded video of a pipe with strong light distortions and color variation. This video was less clear than the vision seen by the sub when tested in Teagle pool, however, it helped the A Team tune values to deal with non ideal vision. The pipes vision is relied on in two parts of the competition: the first pipe recognition and the second pipe recognition. Although both of these tasks performed flawlessly early in the testing period, changes made while working in the simulator resulted in future problems. The A Team was able to get the pipe tasks working during the pool test prior to the competition pool test, but come competition day errors arose in the sub's ability to align in the correct direction.

4.2 Bins Mission

Due to time constraints, the bins mission was primarily tested in the simulator, a digital simulation of Teagle pool created by CUAUV. Although the mission worked very well in the simulator there were several issues when the mission was tested in the pool. While testing the first version of the bins mission, it was discovered that the sub shifted away from the bins while lowering. This issue was resolved by implementing active centering as the sub lowered. Unfortunately, the markers portion of the bins mission was never tested in Teagle pool.

5 Results

5.1 Pipes Vision

After completing the buoy task, the sub successfully located and aligned to the first orange pipe. While driving towards the second pipe, the sub appeared to recognize an object on the pool floor, possibly a pool lane, and veered off course forcing the A Team to kill the mission. Despite making minor changes in the pipe vision, the sub continued to perform inaccurately and the A Team decided to accept the final score.

5.2 Bins Mission

Since the sub never reached the second pipe, it was unable to perform the bins mission. However, the A Team did modify the bins task code to identify the soda can, which was the designated shape for the competition.

6 Conclusion

The A Team's final score was 1000 points. Due to the obscure error in locating the second pipe, the A Team lost to Kathode Ray Tube by 250 points. Overall, the new member training competition was an effective way for myself and other new members to learn how to implement vision modules and mission code for completing tasks in the pool. Although the Teagle pool arena the new member teams faced was significantly easier than the arduous obstacles of the TRANSDEC arena, the tasks and problems faced strongly paralleled those faced by experienced members when implementing vision and missions for more challenging courses.

Looking back, the pipes vision should have been more robust. This would have prevented the sub from recognizing shapes other than the pipe. Furthermore, the A Team should have avoided changing functional code without testing such changes in the pool first. The A Team's primary shortcoming was not utilizing pool tests efficiently. Due to cramming last minute coding into software meetings prior to pool tests, code was often not fully functional or up to par with pool test standards. I have learned the value of every minute afforded at Teagle pool and plan to make sure my code and the code of my teammates is prepared prior to the day of the pool test.

Although I am not fully satisfied with the outcome of the pool test, I feel extremely equipped to tackle future challenges and am excited to contribute to future software that will be utilized in CUAUV's final performance at TRANSDEC.