UM::Autonomy's Sea Serpent

UM::Autonomy 2022 RoboNation RoboBoat Competition Final Paper

UM::Autonomy, University of Michigan College of Engineering, Ann Arbor, MI, USA Submitted May 15, 2022

Ryan Do, Tom Gao, Yi Lin Sim, Asheya Naik, Aayushi Nakum, MQ Nguyen, Saheth Edupuganti, Lucas Mitsch, Frances Truong, Lance Bassett, Rajin Nagpal, Chad Rakoczy, Arjun Anand, Allen Ho, Alex Martin

***Abstract*— This paper describes the development process and competition strategy of Sea Serpent, the University of Michigan's boat submission for the 2022 RoboNation RoboBoat competition.**

Fig. 1. Sea Serpent During Mock Competition

I. INTRODUCTION

UM::Autonomy is the University of Michigan's Autonomous boat engineering build team. This is the 16th year that UM::Autonomy has participated in the RoboNation RoboBoat competition. This is also the first year back to being able to work in person, which was beneficial in many regards, but also came with its own set of challenges. Through this year, the five subteams (electrical, hulls & systems, advanced capabilities, artificial intelligence, and business) came together to learn and grow while making Sea Serpent.

II. COMPETITION STRATEGY

The introduction of new challenges in place of old ones in this year's competition necessitated the establishment of a strong foundation for the accommodation of new hardware and software development as well as thorough sensor integration. This strong foundation, the team decided, must compose of a durable hull design and a fast, small, lightweight system that provided enough room for mounting electronics and hardware while also being largely accessible and modular such that any components could be easily reached and removed for troubleshooting. In order to better utilize our time, UM::Autonomy went against tradition and decided to refurbish one of our boats from a previous season rather than building a boat from scratch - this gave the team an opportunity to spend this time elsewhere, while also being able to work with the boat for hands-on testing and design of hardware for the new challenges.

While the team wanted to attempt most of the challenges in the competition, challenges were prioritized based on their worth in points as well as the complexity and time needed to master each challenge, Looking at the specific tasks our team wanted to prioritize, we wanted to focus on completing Navigation Channel, Avoid the Crowds, Find a Seat at the Show, and Snack Run the most. After that our we would consider the water blast challenge and especially the skeet ball challenges as reach goals.

A. Thrust-to-Weight

In 2019, UM::Autonomy earned 140 points in thrust-weight static judging, with a weight of 50lb and a thrust of 20lbs. This year, the team aims to double the thrust with a six-thruster setup that features two sets of forward facing thrusters. However, due to the reuse of a boat with foam hulls as opposed to lighter carbon-fiber hulls, the expected weight to increase to around 70lbs. Overall, the aim is to earn 140 points in this category again. In Figure 2, you will find the Thrust-To-Weight goal setting graph we created to delineate point totals from thrust and weight points from the point system in the rule book. Here, the purple point represents our goal for 2022.

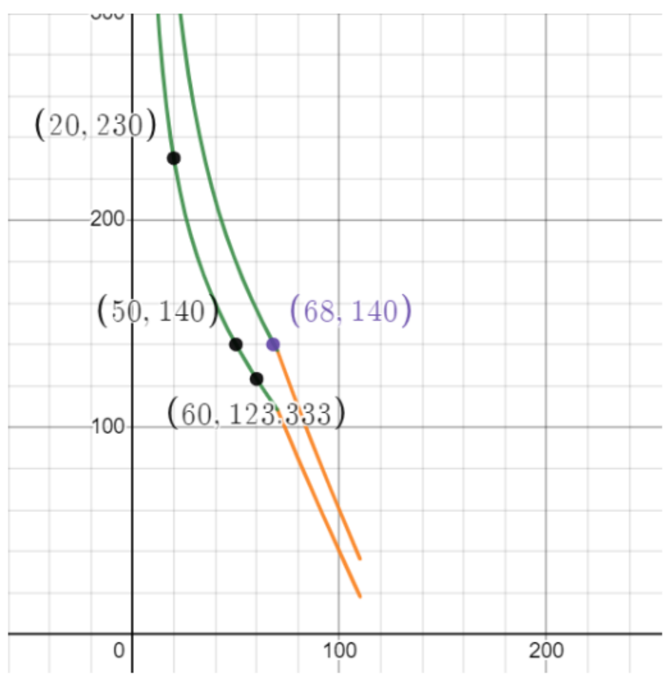

Fig. 2. Thrust-To-Weight Goal Setting Graph

B. Navigation Channel

Given that this challenge is mandatory to attempt any other challenges in the competition, completing this challenge quickly and successfully is one of the biggest priorities of the team this year. While it may not require any fancy algorithms, it still requires all parts of the boat to work together to perceive, localize, and move. In

2019, the boat could only successfully pass the navigation channel once out of four qualification runs as a result of a major electrical failure onboard. As such, the team set this as a High Priority challenge that we aimed to rigorously test prior to competition while also ensuring boat modularity so that any system failures could be assessed separately from the completion of this challenge. The team established a goal to complete the challenge successfully in 14 out of every 15 test runs.

C. Snack Run

As it is a timed challenge with lots of scoring potential that really only requires basic navigational capability and a fast boat, this challenge was a priority for the team this year. After assessing the score breakdown from the 2017 Speed Challenge, it was determined that the team could easily earn the points for entering and exiting the gates and circling the mark, for a total of 250 points. Simply doing this was a High Priority task for the team, with a goal set to complete the challenge successfully at least 9 out of every 10 runs. Based on the 2019 score-sheet, the team assessed that a time between 25-45s is needed to remain competitive in the Snack Run challenge, with the fastest 2019 run coming in at 27 seconds. A baseline goal of 35 seconds was set for the completion of a timed run, with an optimistic goal of 26 seconds.

D. Avoid the Crowds

Since this challenge also requires minimal external hardware or software development and mainly just involves careful navigational operability and fine motor control, this was a High Priority task. One of the biggest advantages that was found is that this challenge could be tested and fine-tuned really early in the development process. The aim was to complete this challenge successfully nine out of every ten test runs.

E. Find a Seat

This challenge is a bit more involved in terms of CV and color/shape recognition, which is why the team declared this a Medium Priority task - still very doable, but something that would come after successful completion and mastery of the Navigation Channel, Snack Run, and Avoid the Crowds. Again, since this did not involve any external hardware development, the team found that it could definitely be accomplished with enough time provided for in-water testing with the boat. For this challenge, UM::Autonomy aimed to complete the challenge correctly in at least eight out of every ten runs.

F. Water Blast

Being one of the new challenges in the competition, the Water Blast challenge was something the team set early on as a Low Priority task. This does not mean that UM::Autonomy did not invest time in it, however - instead, the team repurposed the old Drone sub-team into the Advanced Capabilities sub-team specifically so that the hardware and software required for this challenge and Skeeball could be developed throughout the season. But, since this is the first season with this challenge and hardware and software development of external mechanisms pushed back actual testing, we knew that immediate mastery of this task would be difficult and time consuming, and should only be attempted after other challenges. The team aims to complete the challenge successfully in at least three of every five runs. UM::Autonomy chose to only focus on completing the Water Blast challenge this year, though work was done throughout the year to complete the Skeeball task in the future.

III. DESIGN CREATIVITY

A. Hulls & Systems

After deciding to reuse the Flying Sloth hull (2017) there was much rework to be done. The first step that was taken in the reparation process was to sand off the existing paint to reveal the fiberglass work below. What the team discovered was that the fiberglass and foam structure inside had rotted out in some places and that there were many areas where the fiberglass was bumpy and rough. To fix this the team started by adding a primary layer of Duraglass to reinforce the hull and fill in the cavities as seen below in Figure 3.

Fig. 3. Sea Serpent After Duraglass Installation

Next Bondo resin was applied and sanded to smooth the surface. This process was repeated about seven times until the hull became smooth and reached the desired contour. During this process one and a half gallons of Bondo was used, along with half a gallon of Duraglass. Then surface was sanded with sandpaper up to 1500 grit, as seen in Figure 4 below. Then the hull was ready for painting and waterproofing.

Fig. 4. Sea Serpent Final Resin Sanding

For the painting process the team started off with adding a layer of primer and sanding it. This process was repeated several times until the boat had a thick layer of primer. Next the white paint was applied and the University of Michigan stripe detailing, as seen in Figure 5 below. Finally the boat was coated in a layer of clear coat to add a layer of water resistance and to apply a gloss finish.

Fig. 5. Sea Serpent Post Clear Coat

The last thing the team did was apply scale-like detailing to our aluminum deck using a rotary brush tool, to match the chosen name of the Sea Serpent as seen in Figure 6 below.

Fig. 6. Deck Detailing

When determining how to integrate this older hull with our current thruster setup, the team found that the existing thruster mounts were designed for an older model of thruster that the team had phased out years back. To be able to mount our current thrusters (BlueRobotics T200), a new pair of aluminum rails was fabricated to serve as an interface between the boat and the thrusters. Historically, one of our subteam objectives has been to develop a flexible system where thruster positioning could be adjusted prior to deployment, so that optimal placement could be determined experimentally. The team was able to accomplish this with these thruster rails by drilling as many mounting holes as was possible into the aluminum bar, and devising a system to easily bolt and unbolt thrusters. This multitude of mounting holes is shown in figure 7 during the initial drilling attempt. The flexibility of this system also allows us to easily expand our current 4 thruster setup to, for instance, a 6 thruster setup, if it was deemed necessary.

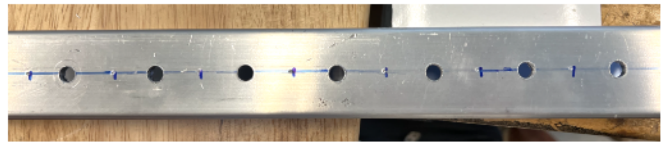

Fig. 7. Thruster Rails at Beginning of Manufacturing

B. Artificial Intelligence

A new method to detect objects such as docks and buoys was implemented using computer vision. Images from previous years were labeled and then used to train a YOLOv4 machine learning algorithm. Training was done using Google Colab to allow for easy and repeatable training using different data sets. The various data sets were made using image augmentations to change the images so the model did not overfit to training data allowing for better performance. With the help of the onboard GPU the model is able to quickly and accurately detect and classify objects as seen in Figure 8 below.

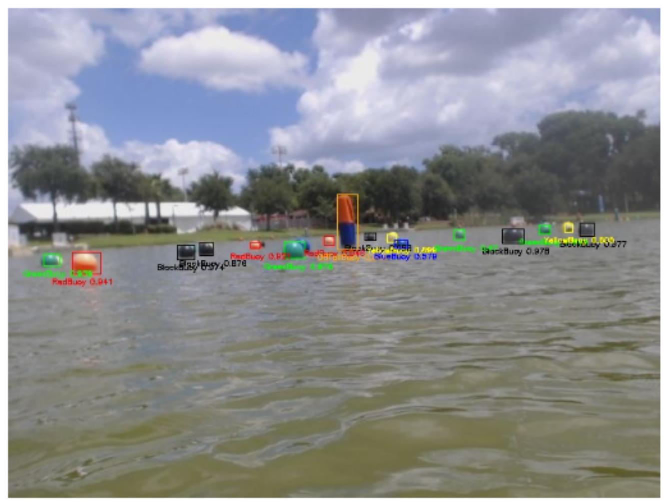

Fig. 8. CV Object Detection

Previous years planning algorithms implemented a Theta* approach for more optimal results compared to an A* based algorithm at the cost of higher compute times. In an effort to be able to run the planning algorithm more often for controls, this years planner switched to a Hybrid A* algorithm that trades faster compute times for less optimal paths. The Hybrid A* algorithm works by creating a tree starting from the start node where any two adjacent nodes are connected via a precomputed motion primitive. The primitives are computed using a linearized dynamics model which accounts for fluid drag. The cost assigned to any given node was the sum of the cost from the start to the node plus the estimated cost from the node to the goal. The estimated cost was calculated by taking the maximum of the Reeds-Shepp cost and the cost from running vanilla 2D A* on the node. Using the maximum of these made the heuristic admissible.

C. Systems Engineering

An overall change regarding team organization involved the use of Jira and Confluence for task tracking and technical documentation. Jira facilitated breakdown of tasks into smaller, more manageable tasks and allowed members to always be up to date regarding progress and also be able to delegate tasking among themselves. This would work together with Confluence, where anyone can add the technical knowledge on any particular topic to be accessed as a reference by others.

IV. EXPERIMENTAL RESULTS

A. Superstructure

This year the team chose to create a superstructure to our vessel. The purpose of this was to protect our sensors, improve the overall look of the boat, and to develop skills in mold manufacturing and carbon fiber layup.

Therefore, in order to be able to make this super structure, there were two test runs of carbon fiber layup attempted - with minor variations from one to the next in order to have a better outcome. The first run did not have a good result, which was concluded to be because there was not a good enough vacuum created. An image of the first run is shown in Figure 9. The second run had a much better result, as the issue with the vacuum was much improved. The set up and final product from the second run is shown in Figure 10.

Fig. 9. Run one of carbon fiber layup

After practicing carbon fiber layup, the team then proceeded with developing the superstructure. To do this the team started by creating a CAD model of the superstructure in Rhino 3D, as shown in Figure 11. From there it was sent to Fusion and created a female mold of our superstructure.

After considering the height limitations of the Shopbot Router that would be used to CNC the mold, the mold was split into three different parts. At this point the team was ready to create the molds as seen below in Figure 12.

There were a few options for what to use as a deck material, either now or for future seasons, so its ability to withstand a load, and be fastened to other surfaces was tested. The main material that was being considered is shown in Figure 13. The main comparison point for this material was a plank of wood - as used in previous years.

To determine if it was worth switching to the honeycomb mesh material, the team ran several tests on its strength and usability.

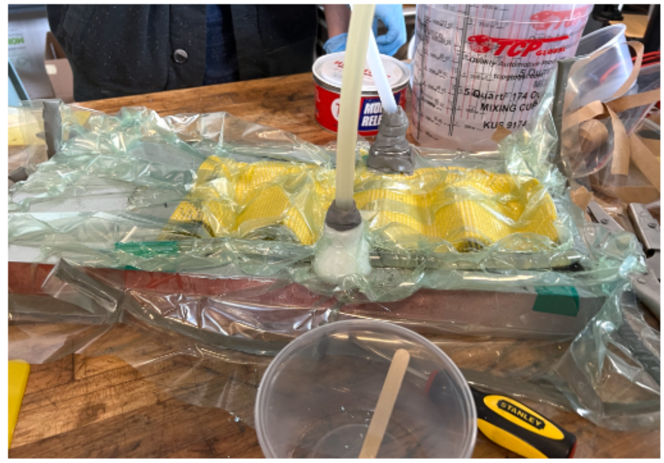

Fig. 10. Run two of the carbon fiber layup, with the result

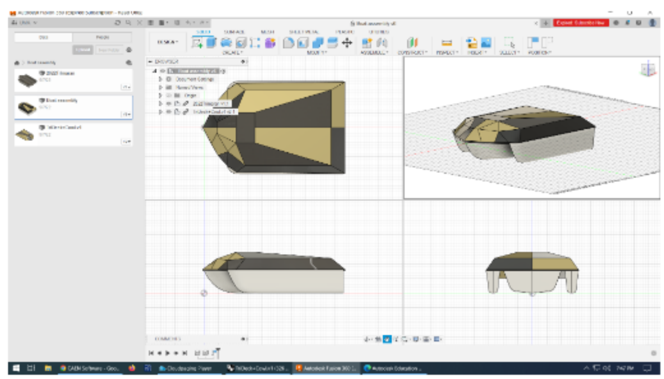

Fig. 11. Fusion Model of Sea Serpent with Superstructure

Fig. 12. Using the Shopbot Router to Create the Superstructure Mold

Since the team would be applying the mounts to the hull on the edges of the deck, the team decided to test the material's ability to maintain strength for these conditions. To do this several size holes were drilled at varying distances from the edges of the sample material to see if the honeycomb mesh would crumple during the drilling process. Next, bolts were fastened and the deformation was measured when tightened on the drilled holes. The final test that was ran was placing the deck on top of two spaced planks of wood

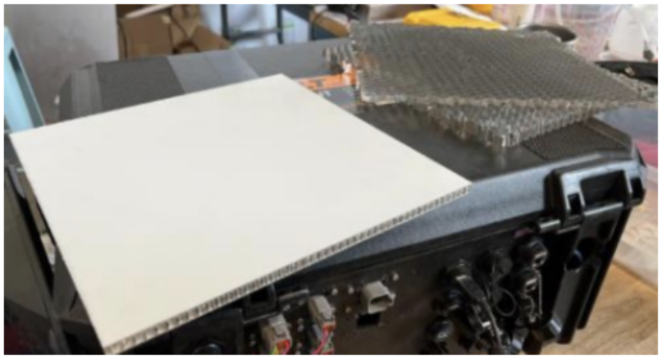

Fig. 13. Potential Deck Material

and applying a heavy load in the center to see if the material would buckle. In the end, the material passed all these tests and greatly surpassed our expectations and mission requirements.

B. LiDAR Deep Learning

For the first half of this year, the AI subteam was exploring the possibility of running a 3D deep learning model for our LiDAR. The plan was to fuse these results with our CV deep learning model's results, providing a more robust dock and buoy detection system.

The team used MMDetection3D's VoteNet implementation as our model architecture, and our dataset consisted of only 529 point clouds from competition ROS bags and data collected on a university pond. The labels included 3 classes: dock, tall buoy, and round buoy.

The team suspected results would be subpar given our limited dataset, and this was quickly confirmed. There was just not sufficient data to learn anything meaningful. The baseline results can be seen below in Figure 14. At best, we achieve a 0.4 mAR at a 0.25 IoU threshold, and our precision at any IoU threshold is very near 0.

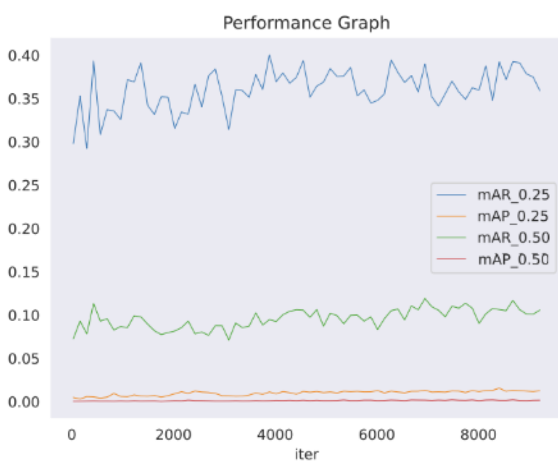

Fig. 14. Baseline results for the team's VoteNet model.

Despite the objectively poor performance, the problems of limited data were approached and class imbalance as one would with any deep learning project. To address class imbalance, we tweaked class weights in the loss functions, where class weight for class X was calculated as (total number of objects / total number of instances of class X). Data augmentations were also implemented to combat our extremely limited dataset. These included point jitter, rotation and scaling, and horizontal flips. The results for mAR and mAP at IoU threshold 0.25 are shown below in Figures 15 and 16. The class weights showed no real improvement, and while the augmentations showed relative improvement, our best results were still well below where they needed to be for deployment.

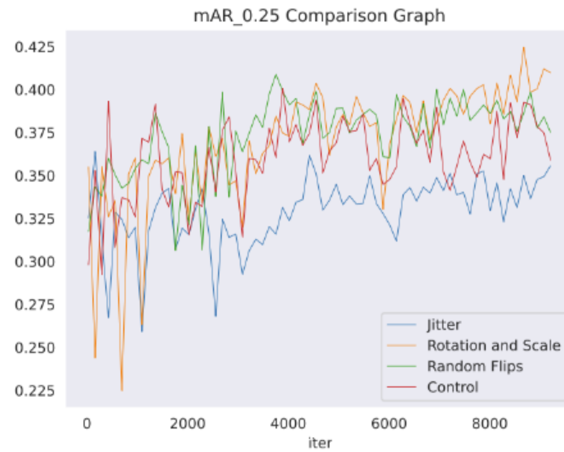

Fig. 15. IoU 0.25 mAR results for various data augmentations.

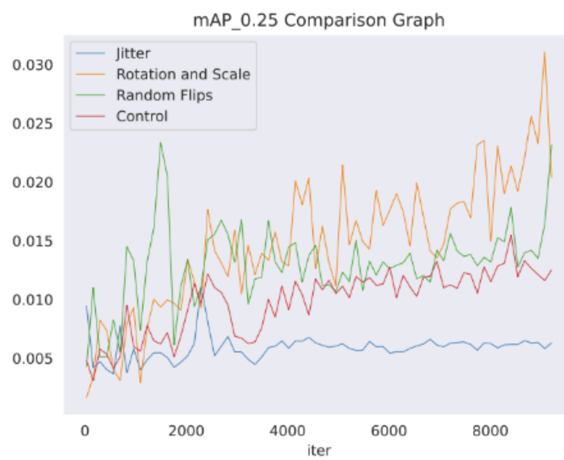

Fig. 16. IoU 0.25 mAP results for various data augmentations.

In the end, the team chose to put this project aside in favor of more competition-focused projects. While interesting, simply a larger dataset was needed for this to be viable and having members spend their semester labeling data was not something that the team was interested in doing, especially given that there are have other methods of detection already implemented. As an experimental project, however, this had a lot of value in teaching members about deep learning workflows and addressing real world dataset issues.

ACKNOWLEDGMENT

We would first like to thank our corporate sponsors for their assistance with fabrication and funding. Special thanks to Ford Motor Company, Boeing, APTIV, and Northrop Grumman for their support. We would especially like to thank Ford for their assistance with mold material and fabrication. Without their help, we would not have the material required to build the super structure of our vessel.

In addition, we would like to thank our university sponsors for their assistance. Without the College of Engineering and its multitude of resources available to student project teams, we would not have been able to compete in the RoboBoat competition.

The existence and success of our team depends on the incredible support of the University of Michigan, our advisor, our committed alumni, and our industry sponsors. A special thanks to Professor Kevin Maki for being the team's advisor and his mentorship in naval architecture.

REFERENCES

[1] Group, Sleipner. A Guide to Marine Bow and Stern Thrusters for Leisure Boats, info.side-power.com/guide-marine-bow-and-stern-thrusters.