INDIVIDUAL LAB REPORT 2

March 7, 2019

John Macdonald

Wholesome Robotics, Team E

Teammates:
Aman Agarwal
Aaditya Saraiya
Hillel Hochsztein
Dung-Han Lee

Contents

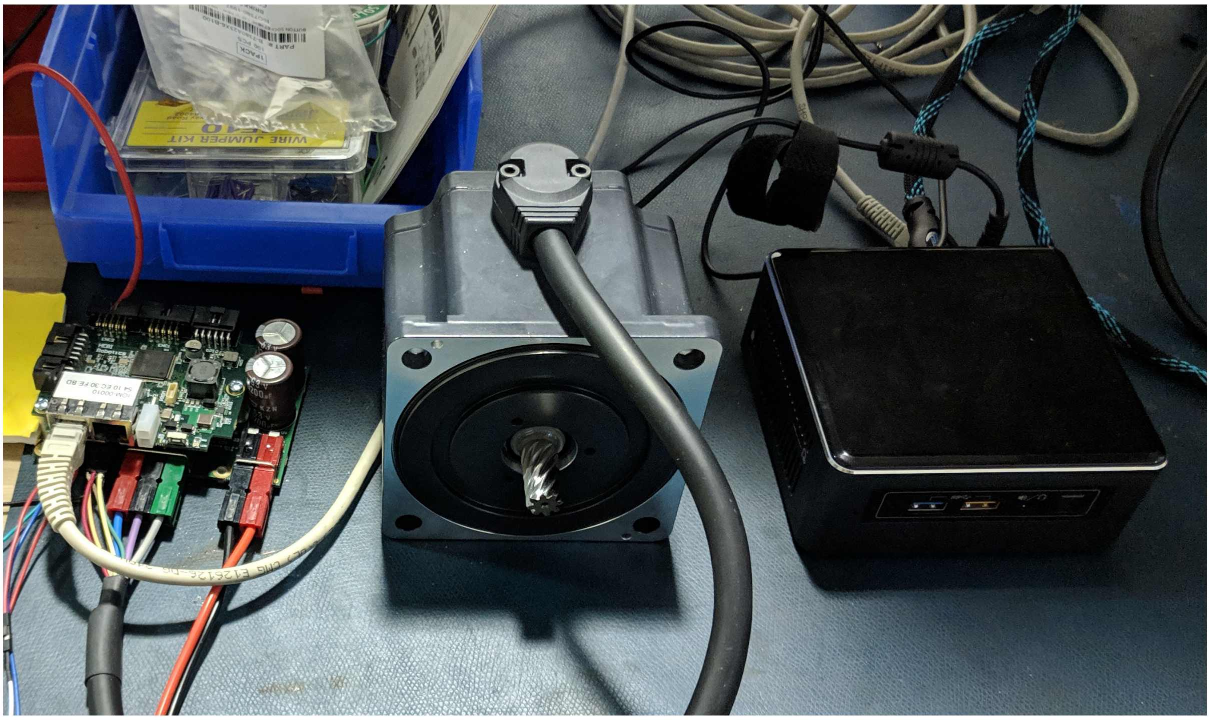

Figure 1: The Intel NUC with the motor controller connected via Ethernet and the motor controller E-stop signal connected to the power supply in order to emulate and E-stop.

0.1 Individual Progress

I focused more this round on the robot platform. I completed the following tasks related to platform integration: (1) I set up the Intel NUC computer with ROS and installed the software needed to control the HEBI motors; (2) I set up a wireless network for our NUC so that I do not need to store my Andrew ID credentials on the robot in order to get internet access; (3) I installed and set up a DHCP sever on the NUC which is required in order for our custom HEBI motor controllers, which are Ethernet devices, to obtain IP addresses; and (4) I set the NUC BIOS to boot when the NUC recieves power, ignoring the power button, which will be obstructed when we construct the robot. With some help from HEBI, I was able to control the motor from the NUC.

On the algorithm side, I did implement a new ROS Node which reads the RTK GPS ground truth data, transforms it into the local row frame, and compares it with the robot pose estimate in the row. The results will be discussed in the Challenges section.

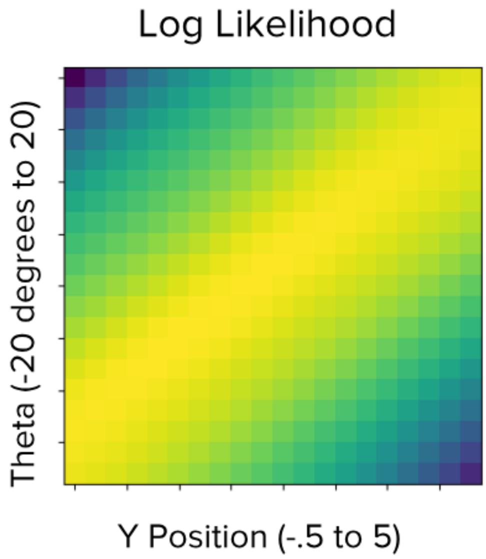

Figure 2: A plot of log likelihood versus angle for a single example sensor reading. Zero angle is facing straight down the row, and y position in meters, the direction moving laterally between rows. It would be expected to have a peak in the center, however, there is a problematic amount of likelihood associated with other poses.

0.2 Challenges

The biggest challenge has been trying to find the right balance between robot platform work and algorithm work. We began increasing our work on the robot platform in Sprint 2, rather than focus on algorithm work as we had in Sprint 1. We anticipated a significant lead time for manufacturing the parts and wanted to minimize the amount of additional time spent on integration once the mechanical parts arrive. As a result, I spent a significant amount of time familiarizing myself with the existing platform's software, especially related to motor control. We have very nice custom motor controllers from HEBI which connect over Ethernet, however because they are custom documentation is very limited. The motor controller requires an active "un-stopped" signal coming from an e-stop in order to respond to any commands, which unfortunately wasn't documented and was discovered when I arranged a meeting with the HEBI representative on-campus to help us with the motor controllers.

At the same time, we've realized that because we are just now requesting a quote from a machine shop for our robot platform parts, given the assumed lead time of about a month (potentially more), we will definitely miss our target goal of having a driveable platform by the end of March. The new thought from Aman is that it is unlikely we will even be able to assemble the robot platform by the end of this semester. This has caused us to re-evaluate whether it makes sense to continue pushing on the electrical and software side, if the mechanical part of the platform is going to be so delayed. This reality has caused some indecision, however, our current direction is to continue to work on the electrical and software integration, since we will have to do that eventually anyways, and speak with George about potentially using the existing Robotanist platform for longer.

With regard to localization, the current Gaussian mixture model has not progressed as much as I would like. The major difficulty with this model arises from the fact that for positions off center of the row, it finds angles at which the projected points "look" like a wider row of two plants. This causes the particle filter to diverge with large angles. It appears that because the likelihood issue is "balanced" along Y, the Y estimate remains relatively accurate. In theory, the distribution is also "balanced" along angle, however this diverges quickly.

The mean error for the filter, collected over about four meters of robot data, before the filter diverges, is: X: 3.3 meters, Y: -0.0052 meters, θ: 1.51 radians. Due to lack of improvement of the model, the current direction is to try a line-fitting approach for the sensor model and see how that compares in accuracy.

0.3 Teamwork

We continue to work on the following domains:

- John: Software, robot localization
- Aaditya: Software, SLAM
- DHL: Software, plant health indicators
- Aman: Hardware, robot platform
- Hillel: Hardware, weeding

As my work this round focused more closely on the platform, I collaborated with Aman to determine suitable mounting locations for the sensors. I also collaborated with Hillel in the mounting of the lidar sensor.

Additionally, I discussed evaluation metrics for the plant health deep net with Dung-Han and possible avenues for improving the accuracy. We are not sure why our model is not detecting health indicators with the same accuracy which Harjatin, a graduate student in George's lab, is achieving with a similar method. We discussed the differences between our approaches and where our model might be lacking. We also discussed an alternate method whereby we might combine deep learning with traditional machine learning in order to cope with our limited dataset.

Some images collected by the robot suffer from poor lighting. Therefore, we also discussed how I might implement a method for using the lidar to judge the height of the plants and adjust the height of the stereo camera in order to capture better images for the plant health model.

0.4 Plans

We have scheduled a farm visit on Tuesday March 12th; a van is rented and we are familiarizing ourselves with the equipment needed. Our goal is to collect data with our new Velodyne lidar (previously we used a Quanergy sensor). This will enable testing of the LOAM SLAM algorithm, whose parameters are currently set up for the Velodyne sensor. Additionally, we have a sensor mount for the Velodyne which can tilt as well as move up and down along the main aluminum extrusion structure. By testing several sensor mounting positions, I hope to find a more optimal one than the current setup. If I am able to implement a line-fitting sensor model in time for the test, I may test this algorithm and obtain preliminary results.