Progress Review 3

Rick Shanor

Team D: Human Assistive Robotic Picker

Teammates: Alex Brinkman, Feroze Naina, Abhishek Bhatia, Lekha Mohan

ILR04

November 13, 2015

Individual Progress

This week, I worked on projects for the computer vision and object manipulation tasks. First, I improved the image recognition pipeline such that the system outputs the approximate item coordinates on the shelf. Next, Abhishek, Lekha, and I worked together to run a 1000 image test on the item recognition pipeline. Finally, I led the PCB layout for the suction system. In addition, I designed an enclosure to house the electrical system.

In order to determine the item position on a shelf, I first improved the clustering algorithm. I used the fact the number of items on a shelf is known in order to merge clusters. For this week, I simply merged clusters based on Euclidian distance between clusters. However, this algorithm will be improved to first merge clusters which overlap in the y or z direction.

After accurate clustering, the vision pipeline next estimates which cluster corresponds to each item. This step is achieved by comparing the cluster frontal surface area with a bounding box fit around the known ground truth model. Two successful identification attempts are shown in figure 1. The green box bounds the item of interest. The x position of the item is determined by averaging the left and right edges of the bounding box.

Figure 1: Successful Item Identification

After setting up an end-to-end vision algorithm, I tested the algorithm on 1000 images provided by Rutgers University. This dataset also includes the ground truth position of the item of interest. I worked with Lekha and Abhishek to write the scripts to automate the testing process. The results of this test are shown in figure 2. The algorithm accurately predicted the location of the item on the shelf to within 3 cm on 56% of attempts. After investigating the results, the pencil cups, the dish brushes, and the highlighter packs were especially difficult to detect. All these items were recognized in under 33% of attempts. In order to improve algorithm accuracy, these items might need to be handled in special cases. By the fall validation experiment, we have committed to and will achieve 70% accuracy on live streamed Kinect data.

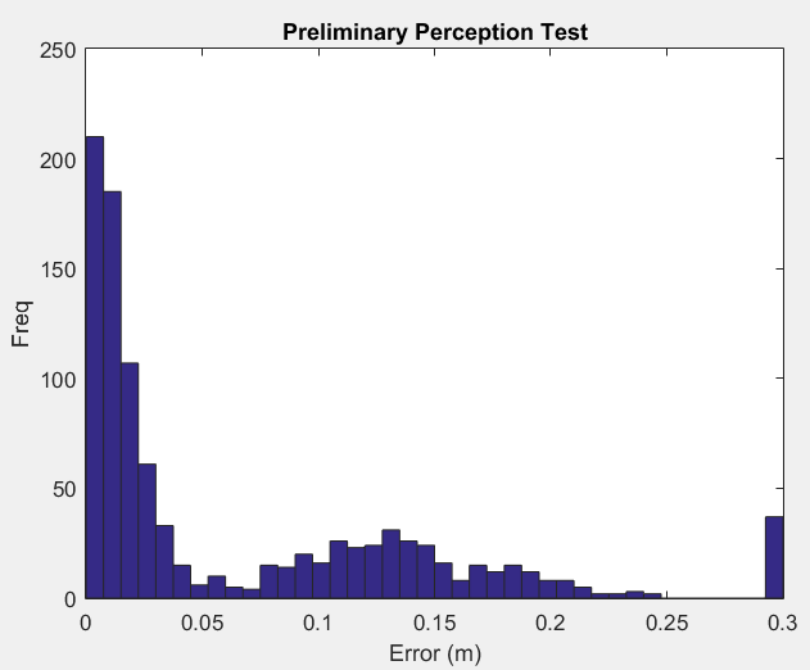

Figure 2: Preliminary Perception Test Results

This week I also designed the gripping electronics system. First, I laid out the components for the PCB and prepared these files for manufacturing. The final board holds an Arduino, reads up to four analog pressure sensors, and controls two AC relays. In addition, I laid out an electronics enclosure that will house all this equipment. As shown in figure 3, the box will have an inlet power connector, two output plugs to connect two vacuums, two serial connectors to attach up to four sensors, and status LED's. This box will be fabricated and the electronics will be assembled next week, assuming the PCB arrives.

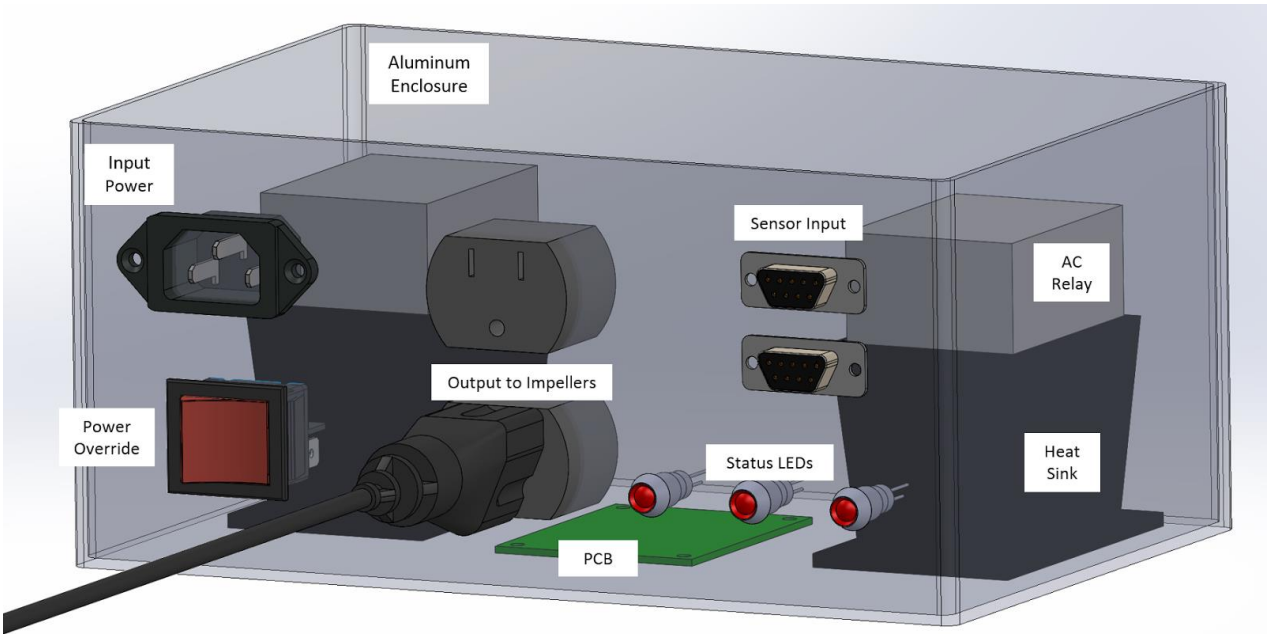

Figure 3: Gripper Electronics Enclosure

Challenges

It took me a long time to converge to an item identification algorithm. First, I tried using 3D feature matching algorithms. However, since many of the items in the APC database are deformable, this algorithm often failed to find useful matches. Next, I tried using iterative closest point methods to minimize least squared cost between the ground truth models and items identified on the shelf. This worked in some cases, but without writing complex algorithms that incorporated knowledge of the cameras perspective to determine a matching score, this method was inconsistent. Finally, by simply comparing the volume of the bounding box of the identified item to the volume of the known item model, identification accuracies of 56% were achieved.

Teamwork

Lekha and Abhishek both assisted in running the perception test. Lekha wrote code to load the ground truth point clouds while Abhishek wrote a script to extract the actual item position from the Rutgers data. Alex led the reiteration of the prototype suction gripper. The new design allows for the vacuum tube to be easily connected. Alex also began attaching the gripper to the PR2 in simulation. Feroze and Alex worked together to get the PR2 arm moving between desired states. Feroze specifically created a ROS service that managed path planning and execution. Alex handled the state controller behavior and TF frames between the markers and the robot.

Plans

Next week, I want to complete the electronics box for the suction system. This includes cutting holes for all the connectors, mounting the components, wiring all the components, and testing the components functionality. In addition, I want to be able to run the perception algorithms on real time data through the Kinect. I want to be able to set up items on a shelf, type the item of interest into the command line, and recognize this items position on the shelf. These two deliverables are critical path in respect to the Fall Validation Experiments.