Progress Review 2

Rick Shanor

Team D: Human Assistive Robotic Picker

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

ILR03

October 30, 2015

Individual Progress

I made progress on both computer vision and object manipulation tasks this week. First, Alex and I worked together to converge upon a gripper prototype that the PR2 could hold and that could pick up at least 7 items from the APC item database. Next, I wrote an algorithm that determines the position of each item on a shelf. Finally, I wrote a script that converted over 10,000 images from the Rutgers dataset into pointclouds. I used (a few of) these to test the clustering code.

After testing several suction solutions, including a Bernoulli vacuum and a DC motor vacuum pump, the shopvac performed the best due to the high air flow and sufficient pressure. The high air flow allows for the suction cup to hold items even when the seal isn't perfect. This week, Alex and I both made CAD models for a prototype. The gripper, which can be seen in figure 1, was designed to be simple, inexpensive, and easy to fabricate. I laser cut an extension which holds the air tube and mounts to a PR2 gripping block made by Alex. I modified an off-the-shelf suction cup to allow for higher air flows and made a mount to interface between the hose and the cup. After completing the prototype, I tested its ability to pick up objects, Figure 1 shows a test of the suction cup to pick up the spark plug. The system also picked up the Oreo container, safety glasses, a sleeve of plastic cups, Elmer's glue, a bag of squishy balls, and two dog toys. A full video can be seen on the team website under documentation.

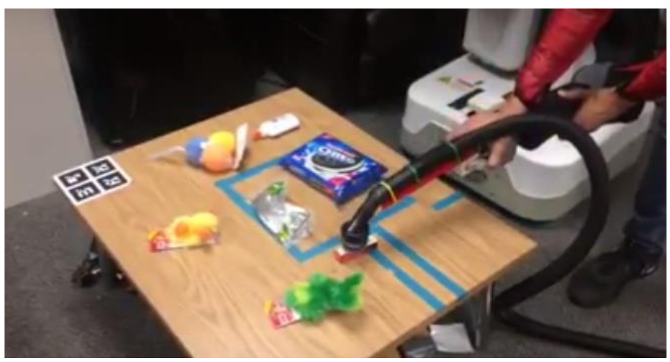

Figure 1: Testing Item Pickup

Since I am still fairly new to 3D object recognition, I read several more papers to better understand state of the art techniques. Lekha and I decided upon the general process shown in figure 2. The first step in item recognition is capturing a Kinect image. Next, we can crop everything out of this image besides the shelf of interest to improve algorithm speed and accuracy. Next, the color image is used to divide the scene into small segments. These segments are clustered based on proximity and expected number of items on the shelf. Finally, matching algorithms are used to determine first which item is the one of interest. Ground truth models are used to estimate item pose. Having established this process, the CV workload is more manageable and distributable across the team. This week, I focused on segmentation and clustering. It is worth noting that green items in figure 2 are working, yellow items are in progress, and red items are yet to be started.

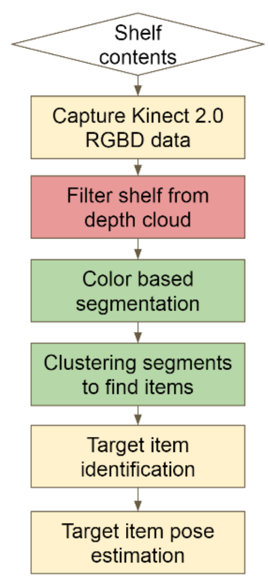

Figure 2: Image Process Pipeline

This week, I implemented segmentation and clustering algorithms. I used naïve methods to filter out the background beyond the shelf simply based on the Z distance measurement. Next, I used Point Cloud Library to segment an image based on color. The result of this process is the blobs of color seen in the output of figures 3 through 5. Using these segments, I estimated item location assuming that all items existed only along the X axis. I averaged each segments minimum and maximum values to determine the centroid. I then sorted the segment centroids from smallest to largest. Then, starting from the leftmost segment, I clustered segments together comparing distance to left and right neighbors. Finally, I used PCL tools to visualize the point cloud and draw bounding boxes around predicted item locations.

Results of these algorithms are shown in figures 3-5. Overall, these algorithms worked very well. In figure 3, all three items are correctly boxed. However, some of the shelf is merged with the Elmer's glue. In figure 4, both items on the shelf are found. The clustering code does fail to merge all the ball clusters and actually clusters them into two items. Finally, figure 5 results are less accurate because of the excessive amount of extra data that wasn't properly filtered out by using naïve methods.

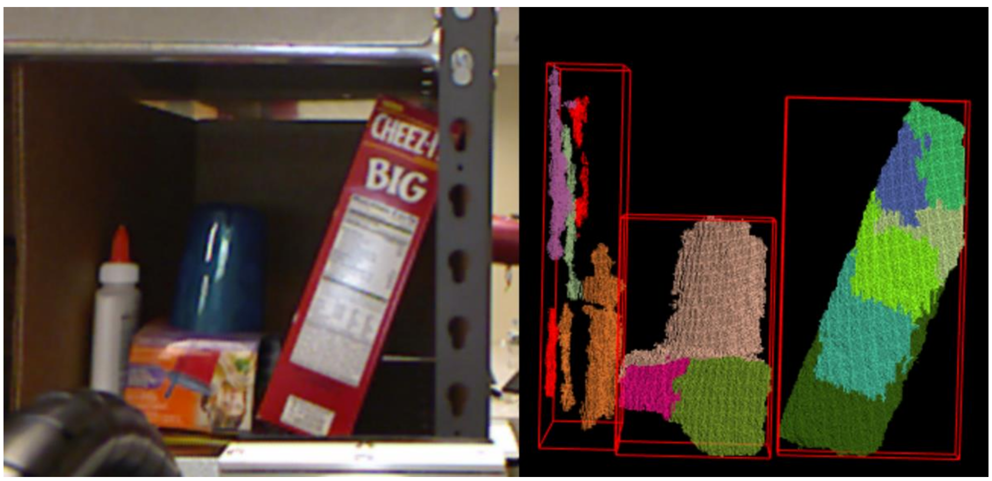

Figure 3: Raw Image (Left) and Output Clustering Results (Right)

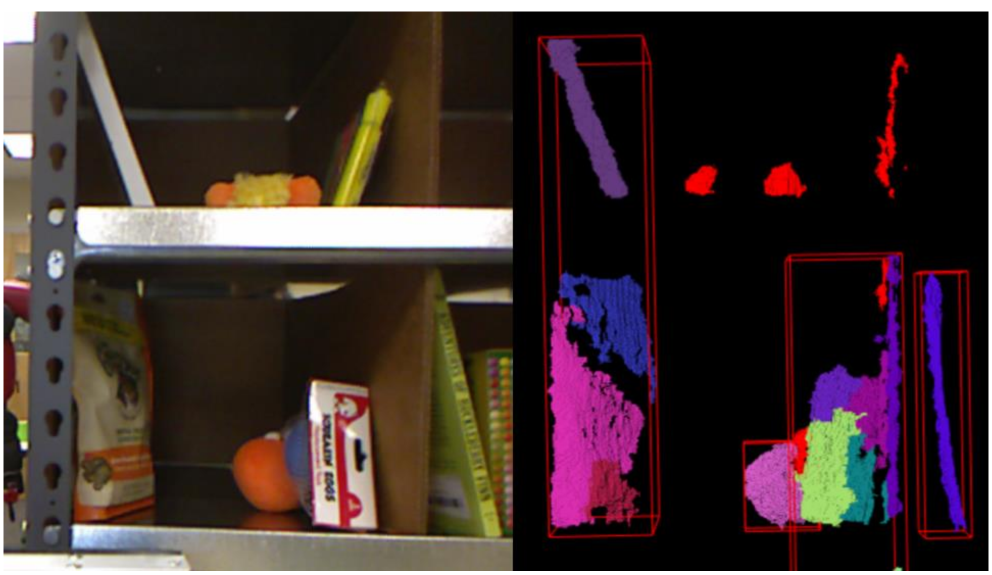

Figure 4: Raw Image (Left) and Output Clustering Results (Right)

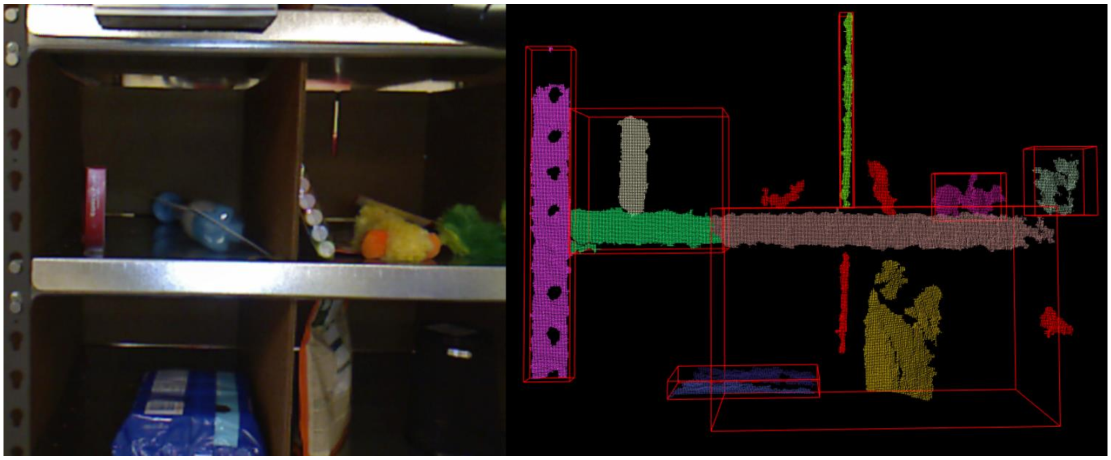

Figure 5: Raw Image (Left) and Output Clustering Results (Right)

In order to improve these results, smarter filtering methods will be used to remove the shelf from the scene before clustering. In addition, because the APC data file tells expected number of items on a shelf, future work will combine clusters to match this expectation, mitigating instances such as in figure 4 where the bag of balls are split.

Challenges

My biggest challenge continues to be my inexperience with C++. I spent several hours debugging my clustering algorithm. For example, I didn't realize that minFloat was different than –maxFloat. Similarly, I didn't understand the cout command might or might not print properly unless manually flushed. As a team, we have had major struggles getting the Kinect 2 to work with our hardware. We have discussed the option of buying a secondary computer to handle this data and pass it on to the main PR2 computer.

Teamwork

This week, Alex continued building the SMACH state machine and got the state viewer working. Alex, Bhatia, and Feroze all worked getting the Kinect 2.0 to work. This was an involved process given that we are running an outdated version of ROS (PR2 runs on groovy) and the Kinect 2.0 requires USB 3.0. Bhatia led the push to test PR2 teleop. He found and installed a package that could control the arms with a joystick. Lekha and Alex both helped him run some tests acquiring items. Feroze was able to use inverse kinematics to control the PR2 arm in RVIZ simulation. Lekha continued to familiarize herself with PCL and she implemented an object detection algorithm. Finally, as mentioned above, Alex and I collaborated on the gripper prototype. He fabricated the gripping block that the PR2 holds.

Plans

Next week, I want to study off-the-shelf impellers. Unfortunately, most components I have found so far have poor documentation. The final goal for the gripper subsystem is to find a drop in motor we can use in place of the shop vac given our initial success in acquiring items.

On the computer vision front, I plan to work on components to flesh out the vision process pipeline outlined in figure 2. Specifically, I want to write an algorithm that filters out the shelf from the point cloud data. This will allow for improved segmentation and clustering.