Progress Review 4

Rick Shanor

Team D: Human Assistive Robotic Picker

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

ILR05

November 25, 2015

Individual Progress

This week, I worked on perception and suction system tasks. For the perception subsystem, I was able to stream images from the Kinect 2 to my computer. I improved the image processing pipeline by implementing iterative closest point to estimate item pose on the shelf. For the suction system, I fabricated and wired the electronics box. I also wrote a ROS node that streams and filters pressure data from the vacuum system to predict if an item has been acquired.

I fabricated and wired the electronics box to control up to two grippers. To fabricate the box, I milled cutouts to mount the connectors for input power, a switch for the AC power system, two outlets to connect the vacuums, 9 pin serial cable connectors for interfacing with pressure sensors, three status LED's, and the USB connector for interfacing with the Arduino. The box is shown in figure 1. After mounting the components in the box, I wired the AC power circuit. Feroze handled soldering the PCB. Alex and I ran a series of tests to make sure we could control both vacuums through ROS.

Figure 1: Electronics Box Assembly

After finishing the box, I wrote a ROS node that filters the pressure data to predict when an item is acquired. Specifically, after the vacuum is turned on, a ring buffer stores the last 50 pressure values. The most recent 25 are compared to the previous 25 to sense when there is a drop in pressure. This indicates that an item has been acquired. In order to develop the filter, data was bagged in ROS, converted into CSV format, and imported into Matlab for easy visualization. The Matlab filter was rewritten in python after I verified the results in Matlab. Figure 2 shows some initial filter results. The top plot shows the noisy pressure data (blue) and a rolling average filter (red). The bottom graph indicates acquisition status. Notice that in the graph, the system only detects picking up items on three of four attempts. Thresholds have been tuned since generating this graph to achieve close to 100% accuracy identifying that an item has been acquired.

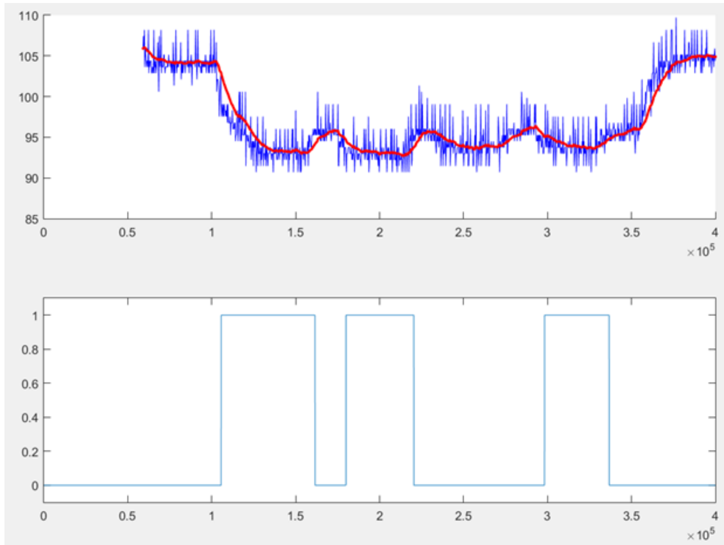

Figure 2: Pressure Data (above) and Item Acquisition Prediction (below)

This week, I started transitioning vision algorithms from running on images in the Rutgers database to running on live Kinect 2 data in preparation for the FVE. I am using libfreenect2 to actually capture the raw Kinect data. Then I am using a converter written by Giacomo Dabisias, a PhD student in Italy, to convert this raw data to point cloud data. I worked with Giacomo to debug an issue that was cropping out depth data under 1 meter away.

Figure 3: Capturing Data with the Kinect 2

Finally, I continued to improve the recognition algorithms. I rewrote my ICP algorithm to perturb the ground truth model such that the solver would not get stuck in local minimum. This allows for better pose estimation of the item in the scene, as shown in figure 4. I improved the recognition algorithm to rely on the fraction of inliers between the ground truth model and the scene model. While this has so far been more accurate than the naïve bounding box approach, I need to further incorporate number of outliers into this scoring metric to achieve accuracies significantly above what we promised for the FVE.

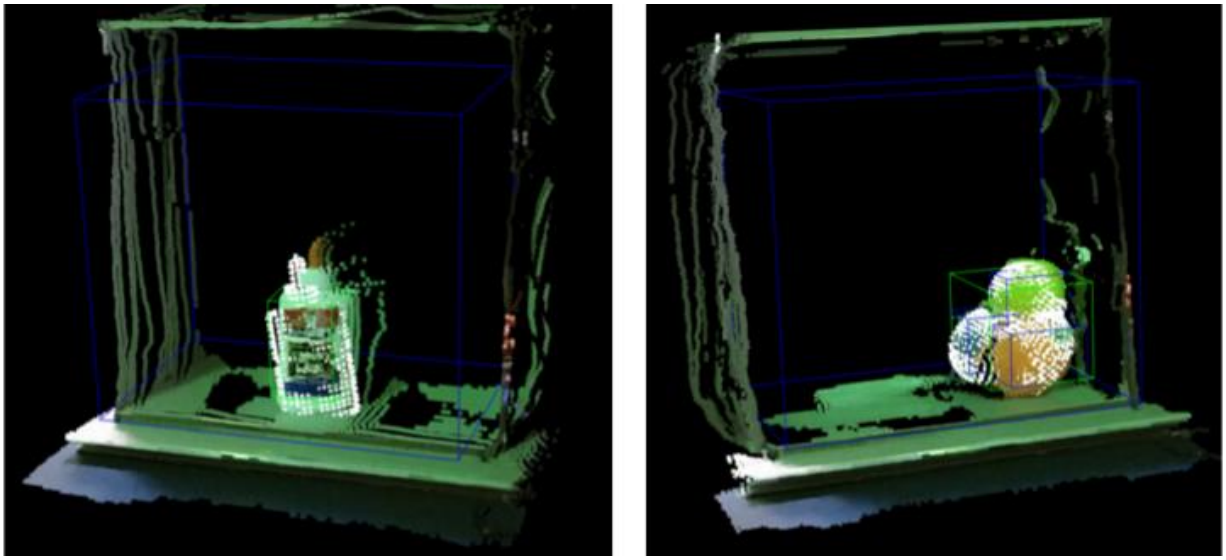

Figure 4: Pose Estimation (white) using Iterative Closest Point

Challenges

The biggest challenge this week involved working out a bug converting the Kinect 2 data from raw format to point cloud depth (pcd). The Kinect 2 is significantly less supported than the Kinect 1 in the development community and I was using a converter I found online on an individual's GitHub site. I worked with the developer to find a bug and get the Kinect reading from .5m away from the shelf. The second challenge has been developing a metric that quantifies the fit of the iterative closest point method for item identification purposes. Simply counting the fraction of inliers matching ground truth to scene has proven too noisy to achieve good recognition.

Teamwork

Lekha has been developing methods to grasp items. Right now, she is simply approaching items from above the centroid, but these methods will get more advanced as we transition to full system tests next semester. In addition, she has been working on interpreting the work order and passing messages in ROS that indicate which item we want to identify. Abhishek has been working to set up the vision computer. He had an initial problem getting the wireless network card working in Linux. Now he is working on a Kinect 2 image publisher in ROS. Feroze soldered the PCB and Alex has been helping with testing the pressure sensor. Alex and Feroze also tried out a current sensor to detect item pickup. Finally, Alex and Feroze have been working on the simulation. They have implemented arm and base control to iteratively move the arm to each shelf bin then move the arm and base to the order bin.

Plans

There are several tasks I need to handle before the Fall Validation Experiment. My main focus will be on improving perception accuracy by developing a better scoring metric to compare the ground truth and scene. In addition, I will finalize the test procedure for setting up items on the shelf. Finally, I will work to transfer these algorithms into the ROS framework.