Progress Review 1

Rick Shanor

Team D: Human Assistive Robotic Picker

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

ILR02

October 23, 2015

Individual Progress

I made progress on three fronts this week. First, I imported PR2 models into the robotics simulators RVIZ and Gazebo to understand geometric constraints. Second, I continued developing the suction cup gripper and prototyped three different vacuum methods. Finally, I worked on implementing item recognition using both 2D images and 3D depth data.

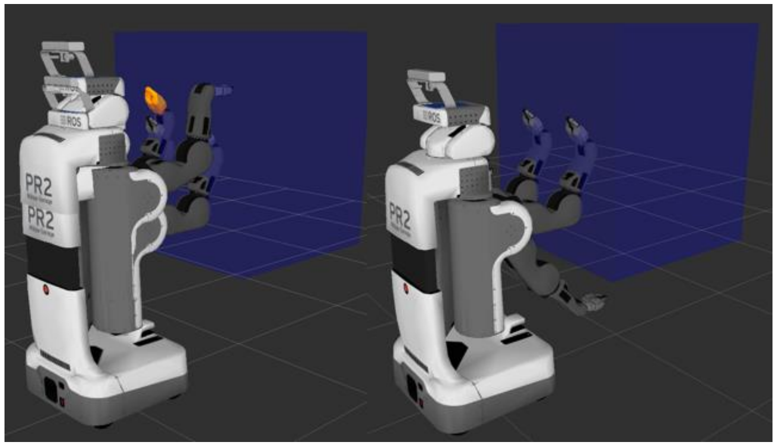

Figure 1: PR2 Workspace Visualization in RVIZ

In order to study the feasability of using PR2 for the Amazon Picking Challenge, I pulled in a PR2 model into Gazebo and RVIZ. Using the MoveIt package, I was able to configure PR2's arms into the maximum (Fig 1, left) and miniumum (Fig 1, right) reach positions. The blue box in figure 1 represents the shelf that the robot must pick item off of. This test showed that in order to pick items off the top shelf, PR2 will need to operate on a platform. In addition, a second Kinect will be necessary in order to see the bottom shelf.

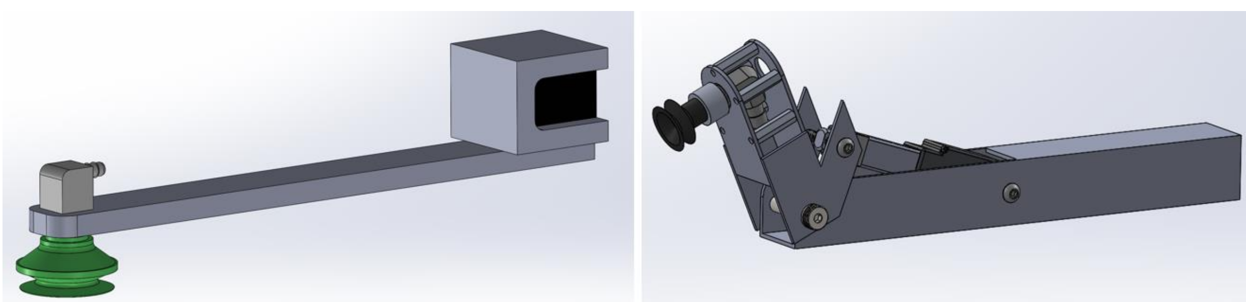

Figure 2: Passive (left) and Active (right) Gripper Concepts

Next, I continued developing the gripping mechanism. Figure 2 shows two different potential concepts for the gripper. The passive mechanism, configured in a cantilever beam configuration designed to bend a few millimeters under pressure, would add compliance and ensure we don't damage objects. The PR2 would hold the block between its fingers so that no physical modifications of the robot are required. The gripper on the right incorporates a linear actuator that actively rotates the suction cup up or down. This would allow for both pressing and pushing actions when acquiring items. Ongoing trade studies and prototyping will help the team choose a concept in the next two weeks.

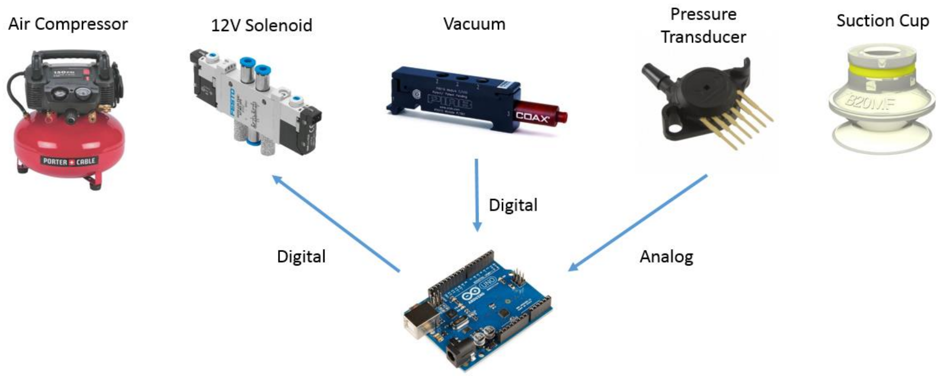

Figure 3: Suction System Components

This week, I prototyped several methods of acquiring items, including a vacuum pump, a Shopvac impeller, and a Bernoulli vacuum. While the vacuum pump was small and efficient, it did not generate enough airflow to pick up porous objects. The Shopvac does a good job of acquiring objects without a perfect seal, but can drop objects under dynamic conditions. The Bernoulli vacuum system, shown in figure 3, showed the ability to pick up and hold a wide range of objects. However, I am still researching Bernoulli vacuum systems to select a pump with the right balance of pressure and airflow.

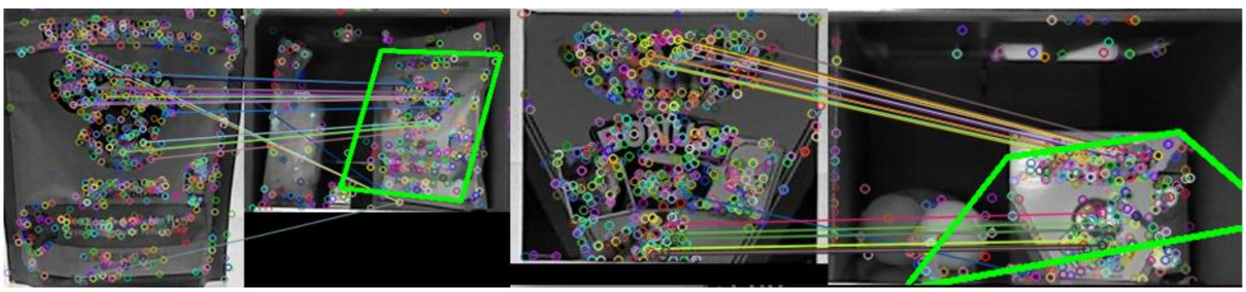

Figure 4: Item Recognition using SIFT Descriptors

Finally, I began prototyping computer vision algorithms to recognize item position and orientation on the shelf. Using ground truth models created by Berkeley and test images published by Rutgers, I implemented both 2D recognition using SIFT and 3D recognition using SHOT. I tested 2D recognition using the OpenCV package. Specifically, I used SIFT to detect and describe features in each image. Then using a brute force algorithm, I matched features in corresponding images. Finally, I calculated a homography transformation using RANSAC to determine an items pose on the shelf (figure 4, outlined in green). Overall, this approach was able to detect items with reasonable accuracy but failed on occluded objects and objects lying flat on the shelf.

Next, I used the Point Cloud Library to estimate item pose based on depth data. Specifically, I used Meshlab to convert ground truth .STL models to a point cloud. In addition, I wrote a script that could convert Kinect depth map and color data to a point cloud. Using this data, I was able to determine the pose of an object in world coordinates. Figure 5 shows the Kinect scene (top grey) and the ground truth Oreo model (top yellow). Using SHOT feature description and a FLANN matcher, I was able to accurately predict the location of the Oreos (blue) overlaid on the world (green).

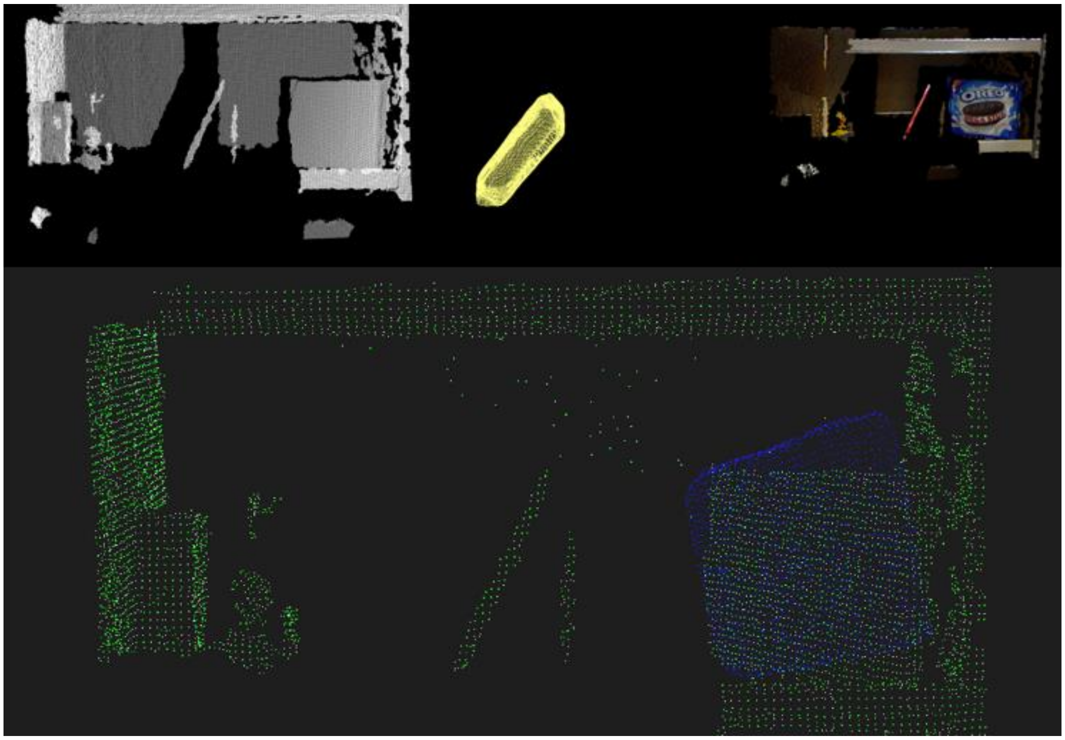

Figure 5: Raw Kinect Data and Oreo Ground Truth Matched to the Scene

Challenges

This week, my biggest challenge was getting the 3D item recognition running. Specifically, I had to go through intensive processes to convert online datasets into a format I could easily work with. Now that this pipeline is set up, I will be able to run more rigorous tests on the vision algorithms. In addition, I struggled learning the basics of C++, the primary language of both OpenCV and Point Cloud Library. This will continue to be a challenge in the coming weeks. Finally, as a team, we continue to struggle with the robotic platform. Specifically, the PR2 has many negatives; it is undersized, overly complex for the task, and moves relatively slow. However, at this point, we don't have an alternative.

Teamwork

This week, Feroze and Alex set up our computer, a team git workflow, and an overall software structure. Feroze wrote skeleton code for the system state machine and set up preliminary ROS nodes. Alex set up a ROS node that communicated with an Arduino and displayed suction pressure. In addition, Alex got the Kinect2 communicating with ROS. Abhishek and Lekha worked on preparing a presentation for Professor Likhachev that highlighted last year's APC competitors. Lekha worked with me on the item recognition front. She investigated ORB and tested algorithms using the VLFeat library.

Plans

Next week, I want to test 3D pose recognition on a large scale. This involves batch processing the Rutgers dataset and running them through SHOT. Also, I want to get a better understanding of how different algorithm parameters effect item recognition. In addition, I want to work with Alex to get the Kinect recognizing objects in real time. Finally, I need to study the tradeoffs between an impeller and Bernoulli vacuum further. By next week, I want to finalize on suction system components.