Alex Brinkman

Team D: Project HARP (Human Assistive Robotic Picker)

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

ILR #: Progress Report

Submitted 3/17/2015

1. Individual progress

I contributed most of my works towards getting the single bin autonomous test working. I added a Cartesian motion planner to improve the path planning inside the shelf, improved the arm control and system architecture, and integrated SBPL onto the UR5. Additionally, I set up a testing rig for a turntable we will use for image database collection out of some spare parts from my personal collection.

The motivation behind the Cartesian motion planner is to improve the trajectory of the end effector inside the shelf bin near the items. The single query planners like RRT or RRT* were often forcing the item down into the shelf or knocking the item against a shelf wall. The Moveit API has a way of computing Cartesian trajectories and has some great qualities. The planning for Cartesian paths is fast – I believe it is using the IKFast solver to compute the Cartesian steps. Another great property is that the function returns how much of the path it can complete and you can plan up until a collision or joint limit, much unlike any single query planner.

Next, I focused on improving the executive scripting of the robot to be more fault tolerant and utilize better target poses throughout the stateflow. If an item is not grasped, the logic now resets and moves on to the next bin, saving a lot of time. Also, I transitioned from using pose set points that are relative to the shelf to set points that are relative to the world frame. This prevents the logic from requesting set points that are outside the configuration space, which have the potential to stop the robot dead in its tracks. This also enables us to precompute trajectories between set points which will greatly improve runtime and guarantee we will move items in a stable, controlled manner.

To enable better quality single query motion plans, I integrated the SBPL planner to run on the UR5 hardware. I encountered many of the same problems when integrating the default configuration and our custom configuration and was able to work through this quickly. The SBPL planner currently does not support the addition of our custom end effector within the planning group, so I had to handle the offset between the UR5 wrist 3 link and our custom suction cup link manually using the TF tree. Additionally, I added the collision model for the end effector as an attached collision object and am currently working through issues with this functionality.

Finally, I pulled together some hardware and actuators for the turntable that will be used to create the database to train the convolution neural net. I used a linear actuator and dynamixel to control the angle of the camera and the angle of the turntable respectively. Rosserial is used to interface with an Arduino that interacts with a high amperage H-bridge to drive the linear actuator with potentiometer feedback to control position. I cobbled a simple fixturing from spare wood from the wood shop over spring break. These components, electronics, and controls come from some past hobby projects and took little time to setup.

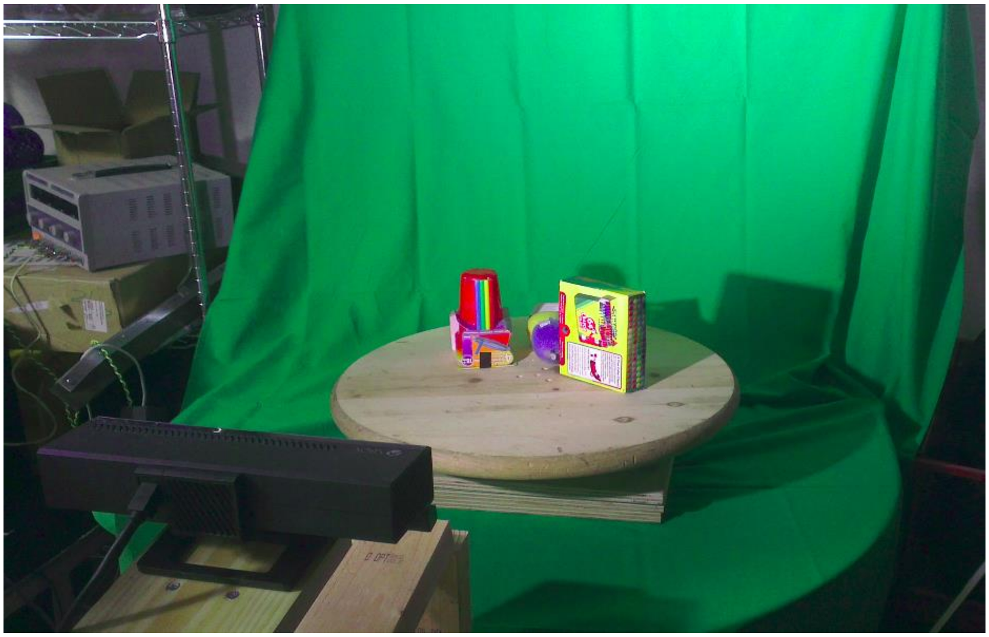

Figure 1: turntable for Automated Image Capture

2. Challenges

The main challenges we faced were with the combined performance of the integrated system. Errors in camera extrinsic calibration, grasping performance, and pressure feedback contributed to unacceptable system performance and will be redressed in the coming weeks.

We estimate there is currently about 1 – 2 inches of deviation from our localized world model of the shelf and our end effector. The most likely source of this error is the camera calibration. To improve this, we attempted to recalibrate the camera intrinsic calibration using the calibration package in ROS but this time we only calibrated the camera and depth sensors within 3 to 6 feet of the camera, reflecting the region of interested in our application. The results showed small but measurable differences in intrinsic parameters and significantly improved our accuracy. However, the accuracy is still only within 2 inches and needs to be improved this week.

Another short coming was the grasp planner. We originally thought the model based approach to estimating the pose of the item in the world would yield superior grasping results to mesh estimates of items but have come to the opposite conclusion. If the ICP algorithm places the ground truth in an incorrect configuration, we may generate grasping points at locations where we do not have any visual confirmation there is an object – in fact, we have data to prove the opposite! This in addition to the deviation of the system leads to low grasp success. We are currently developing methods to do online normal estimation to improve our grasp success once we find the partial set of pointcloud points of the target item.

Finally, the suction system is not working as reliably as it once was. The current system is prone to false positives or false negatives for several different reasons. If the system only partially grasps an item without getting a good seal, the system pressure does not drop fast enough to trigger grasp detection. Sometimes items are detected when the vacuum first turns on despite attempts to tune out this effect. We will redress this issue by improved filtering and are considering various other ways to confirm grasp status.

The final issue I encountered this week is the attached collision object need by SBPL is not working 100% of the time, leading to internal collisions with the end effector. After confirming this issue with Andrew from the SBPL lab, we suspect the problem is in my implementation of the attached collision object interface in Moveit. Work is planned to fix and test this functionality this week.

3. Teamwork

The whole team has settled into the lab and we all spend a lot of our time in close proximity. Rick has been manning the new computer to build up the neural network we think will solve the perception problem of occlusions. Bhatia is helping with perception tools, like masking and segmenting item templates for training the neural net. Feroze was been focused on implementing and integrating the online grasping module. And Lekha is working on designing the neural net and studying prior art in this area.

4. Plans

My plans for the next progress review will be a logical progression of the work I've been doing over the last few weeks. My main tasks are the following:

- To revisit Kinect calibration to improve the deviation of the full system
- Fix the Attached Collision Objects required by SBPL
- Add the ability to precompute trajectories and replay them at runtime
- Finish the turntable to autonomously collect and segment images
- Improve the grasp feedback for the system
- Conduct Failure-mode effects analysis of each subsystem to improve robustness