Feroze Naina M Dheen Mohamed Ismail
Team D - HARP
Teammates – Alex Brinkman, Rick Shanor, Abhishek Bhatia, Lekha Mohan
ILR09
Mar. 17, 2016

Individual Progress

Our system currently estimates the grasping surface of an item on the shelf by fitting ground-truth models and using pre-computed surface normals. This has poor accuracy as it depends on the assumption that the position and orientation of the items are correctly known. However, when working with partially occluded items, we cannot accurately determine the orientation of items on the shelf for ground-truth model fitting.

My solution was to estimate grasping surface in real-time from a partial pointcloud. The points are clustered into grasp surface patches and their normals are estimated. The grasp patches are then ranked based on their quality factor which is computed by finding the dot product of their surface normals with the vertical axis.

This approach was tested using a manually segmented pointcloud of a CheezIt box placed on the shelf. The pointcloud was captured from a Kinect mounted on the arm to simulate operating conditions.

Figure 1 shows the segmented pointcloud and the pose of the Kinect. We can observe that the pointcloud of the rectangular CheezIt box is warped badly. This is because the Kinect is mounted at an angle such that the IR rays projected from the Kinect onto the top surface of the CheezIt box are reflected away from the depth sensor.

Figure 1: Segmented pointcloud of the CheezIt box in the shelf

Figure 2 shows the computed grasp points visualized in RViz. The origins of the red arrows depict the grasp contact point and the direction of the arrow represents the orientation of the suction gripper. We observe that the majority of the good grasping points are on the top surface of the box with the gripper facing down.

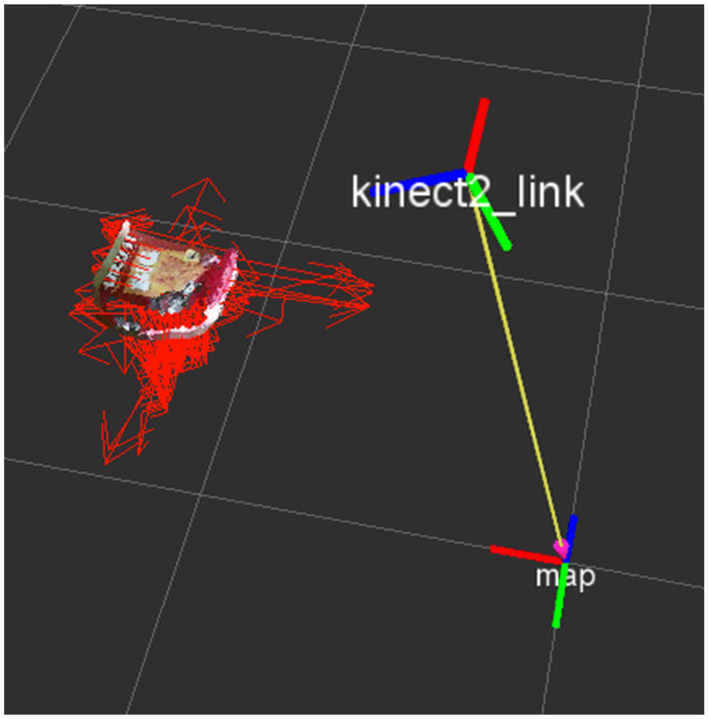

Figure 2: Visualization of generated grasp points in RViz

I am currently working on integrating this with the planning pipeline and testing it on the UR5 robot. The online grasp planner returns a set of grasp poses in decreasing order of grasping quality. The executive planner will iterate through the grasp poses while checking for collisions. When a collision-free grasp pose is found, the grasping is executed.

I also assembled a new computer to run Caffe with GPU acceleration (CUDA) for the perception system.

For the previous project review, I was working on the SimTrack package which uses 3D models and texture for object tracking. However, I could not get satisfactory results and a decision was made to develop our own perception algorithms for handling objects with partial occlusion.

Challenges

It was hard to come up with a good grasping metric for our suction gripper as multiple grasp surfaces would work correctly. If two objects are placed very close with overlapping surfaces, it is difficult to ensure we pick the correct items using the current quality metric. To overcome this, I am exploring centroid-based grasping metrics.

It is also difficult to quantify the effectiveness of the grasp planner due to accumulation of errors in depth camera calibration, localization and pressure sensing. Alex is working towards fixing these issues.

Team Work

Alex worked on integrating Cartesian and SBPL path planners to make arm movements faster and smoother. He also assembled a turntable setup for automatically capturing images and calibrating Kinect for close distances.

Rick, Lekha and Abhishek are working together to test and improve the perception pipeline for multiple objects with occlusion. Rick and Lekha are working on developing a neural network to identify and segment out objects. They have been testing with SegNets, deep learning and the Caffe framework. Abhishek is working to improve the perception pipeline by developing item segmentation masks for use with the neural networks.

Plans

For the next progress review, I will complete integrating the online grasp planner into the system and test it for different objects. I will also be testing a centroid-based grasping approach for convex pointclouds. This would allow us to grasp objects closer to the center of mass and improve grasp quality.

I will be working on collecting datasets for items with different poses using the turntable and segmenting out the items from the green screen. I will also be helping Alex to cache trajectories for increasing planning and execution speed of frequently used arm movements between two points (such as every shelf to the order bin).