



11/13/2015

Progress Review 3

Individual Lab Report #4

Abhishek Bhatia

Team D: Team HARP (Human
Assistive Robotic Picker)

Teammates: Alex Brinkman, Feroze
Naina, Lekha Mohan, Rick Shanor

I. Individual Progress

For the Progress Review 3, I started with working on PR2 base navigation and planning. I spent time understanding the 2D navigation stack available for ROS, specifically the `move_base` package. The 2D navigation stack takes in information from odometry, sensor streams, and a goal pose and outputs safe velocity commands that are sent to a mobile base [1]. The `move_base` package provides an implementation of an action that, given a goal in the world, will attempt to reach it with a mobile base. The `move_base` node links together a global and local planner to accomplish its global navigation task [2]. I spent time on trying `pr2_tutorials` based on navigation stack. But before I could start working on the actual implementation of PR2 base navigation and integrate it with our state controller, I had to shift gears from working on navigation to working on perception subsystem. We as a team analysed our deliverables and felt that we needed to do some load-balancing to help Rick and Lekha with the high priority perception deliverables. We have committed a fairly robust perception subsystem for the fall validation experiment and hence have strict weekly deadlines where we have no room for slips. After analysing these things, we pushed the low priority base navigation by a week and I started working on some of the perception tasks.

The first thing we did was to analyse perception tasks based on the weekly deliverables. Rick was trying Euclidean clustering based image segmentation algorithm (as opposed to color based segmentation where we were not getting great results, Venkat suggested this as a better approach) and the next step in the pipeline was to have a correspondence grouping algorithm to determine the pose of the segmented point cloud by comparing it to the ground truth point cloud. So, we decided to let Rick focus on the image segmentation algorithm and I started working on the correspondence grouping algorithm based on the point cloud library (PCL). I had no prior background in computer vision and so I spent time trying to understand the task at hand, setting up the workspace and getting up to speed as soon as possible [3]. I followed the 3D object recognition based on Correspondence Grouping tutorial on the PCL documentation and made it work on the sample dataset they had provided [4]. Working on this tutorial ensured 2 things, I had my workspace setup and I understood the basic steps involved in this algorithm and got myself familiarized to the point cloud library. The outputs of various pipeline stages of running the Correspondence Grouping on the sample dataset looked similar to the figures below (Figure 1-3).

After this I started working on using the same pipeline on some images from APC dataset. Rick provided me with an image of segmented 'oreo' from the shelf and an image of 'oreo' Ground Truth to test this algorithm. I tried out various configurations of this pipeline, changed algorithms between Hough (a clustering algorithm based on a 3D Hough voting scheme) [5] and GC (a geometric consistency clustering algorithm enforcing simple geometric constraints between pairs of correspondences) [6] and changed various parameters to get some good results but to no avail. The problem that we were facing was that the number of correspondences that this pipeline was calculating on the segmented image was too less than the number required to match instance to the ground truth model. After this I spent time reading about the 2 algorithms implemented in the pipeline (Hough/GC) to understand where we were lagging. Meanwhile Rick was done with the image segmentation task and he took over this task to try out other different things for 3D object recognition and pose estimation.



Figure 1: Sample point cloud with item of interest highlighted in red

(Image courtesy:

http://www.pointclouds.org/documentation/tutorials/correspondence_grouping.php)




Figure 2: Key Point Extraction

(Image courtesy:

http://www.pointclouds.org/documentation/tutorials/correspondence_grouping.php)




Figure 3: Correspondence Grouping and Matching with the Ground Truth

(Image courtesy:

http://www.pointclouds.org/documentation/tutorials/correspondence_grouping.php)

Next, I worked on generating the ground truth point clouds dataset from the Rutgers APC dataset. This was a mundane task as there was no direct method to convert the '.yaml' files that we received from Rutgers dataset to point clouds. We found a script online to convert '.xyz' files to '.pcd' but the only way to convert a '.yaml' file to '.xyz' file was using a linux tool 'meshlab'. I manually converted all the '.yaml' files corresponding to the ground truth of different items to '.xyz' and then wrote a simple bash script to execute 'xyz_to_pcd' conversion script on all the '.xyz' files generated.

Finally, I worked on writing a script to read the '.yaml' file and extract out the 'x' position of the object with respect to the camera from the ground truth image. This was the final part of our perception pipeline where we compare the estimated 'x' position of the object on the shelf bin from the vision pipeline and compare it with the 'x' position of the object in ground truth. Later Rick and myself sat together and integrated everything as part of the perception pipeline and generated the accuracy of 'x' position estimation on a 1000 image dataset.

Besides the work on the perception subsystem, I also worked with Rick on our final board design for the gripper subsystem.

II. Challenges

The major challenge I faced this week was the transition from my work on the PR2 navigation planning and get up to speed on perception tasks as quickly as possible. I received great support from Rick who helped me by pointing me to appropriate resources available online and also helped me understand the vision pipeline. One challenge I faced with respect to the point cloud tutorials was that they contained insufficient information and assumed the user to have some background, which in general if you follow a structured approach should be fine, but for situations with time crunch can be an issue.

Second challenge I faced this week was when I was working on parsing the '.yaml' file of the ground truth to read the x value of the object with respect to the camera. I started working on developing a perl based script to accomplish this task owing to my prior experience in working with perl. Luckily, I found a perl based module to read the '.yaml' files into perl data structures directly, but I failed to make this work after many desperate attempts. Later, I figured out that the '.yaml' format has different ways of representation and the files that we had were not appropriate for this module. So, I worked again to write the script from the scratch treating '.yaml' file as a text and parsed it line by line to extract the information required.

III. Teamwork

The last two weeks were really crucial for us because we had assigned difficult deliverables for ourselves, but our approach to constantly evaluate our status, help each other by load-balancing tasks and prioritizing, dividing and allocating tasks in a strategical manner ensured that we met all our deliverables successfully.

Alex: Alex continued with his work on the state controller and worked with Feroze to integrate arm planning with state controller. He worked on defining the transforms necessary to use interactive markers in simulation and integrated the server provided by Feroze to control arm movement in simulation using interactive markers.

Feroze: Feroze also continued on his task to control arm in simulation using the 'moveIt!' planner. He also worked on generating the server that received commands from the state controller and interacted with 'moveIt!' planner.

Lekha: Lekha worked on with Rick on image segmentation tasks and also focussed some of her time on learning C++.

Rick: Rick primarily focussed on getting the complete perception pipeline in place and generate an end to end accuracy on the Rutgers's dataset.

Abhishek: I worked closely with Rick on perception deliverables. Besides, we also worked together to finalize our suction board design.

IV. Future Plans

This week I plan to get back to my original task of base navigation and planning. Last week Alex started with this task and was able to move the base from one point to another using the navigation stack. I will work with Alex and modify the existing infrastructure to carry out this operation programmatically as part of the state controller. Besides, I also plan to work on defining the shelf transforms in simulation for the demo next week.

IV. References

- 1) <http://wiki.ros.org/navigation?distro=jade>
- 2) http://wiki.ros.org/move_base?distro=jade
- 3) [http://robotica.unileon.es/mediawiki/index.php/PCL/OpenNI tutorial 0: The very basics](http://robotica.unileon.es/mediawiki/index.php/PCL/OpenNI_tutorial_0:_The_very_basics)
- 4) http://www.pointclouds.org/documentation/tutorials/correspondence_grouping.php
- 5) F. Tombari and L. Di Stefano: "Object recognition in 3D scenes with occlusions and clutter by Hough voting", 4th Pacific-Rim Symposium on Image and Video Technology, 2010
- 6) H. Chen and B. Bhanu: "3D free-form object recognition in range images using local surface patches", Pattern Recognition Letters, vol. 28, no. 10, pp. 1252-1262, 2007