Progress Review 10

Rick Shanor

Team D: Human Assistive Robotic Picker

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

ILR09

March 31, 2016

Individual Progress

Over the past two weeks, I have worked on three major perception tasks. We finally received the full 2016 item list from Amazon, shown in Appendix A, and so we can strategize for this year's items. First, I integrated the shelf filtering neural network into our current vision pipeline. Second, I integrated SBPL's PERCH geometric identification into our vision pipeline. Finally, I started data collection and basic testing for our CNN identification approach.

Last week, I demonstrated a neural network that could segment between shelf and items to pre-filter the image and point cloud for initial processing. This week, I wrapped that process up into a ROS package. First, I wrote a program that masks the shelf image, shown in the bottom left of figure 1. This algorithm works by first calculating the coordinates of the shelf corners in 3D space using shelf localization and geometry. Next, I used a search tree to find the closest points in the 3D point cloud to these anticipated corners. Finally, I converted these point cloud points back into image space and applied a polygon mask.

I then wrote a ROS service to wrap around the SegNet CNN I trained last progress review. This ROS node subscribes to the masked image. When the main state controller calls this service, the images is passed into the CNN. The output of the CNN is shown in the top left of figure 1. After the network generates this mask, the mask is then applied to the point cloud. The results of the overall pipeline is shown in the right of figure 1. Note that there is still a small bit of error. This week, we plan to collect more training images with the new APC items in order to achieve even better results.

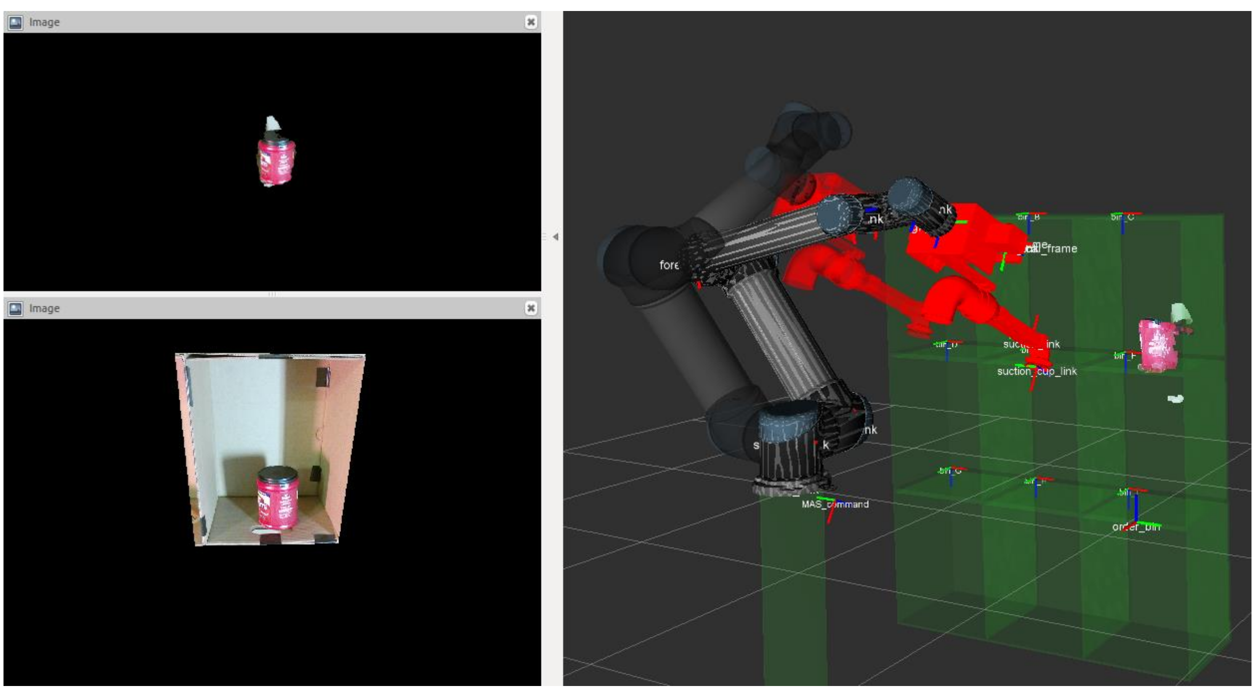

Figure 1: Shelf Filtering Pipeline. Shelf Mask (bottom left), SegNet Output (top left), and Filtered Point Cloud (right)

In addition, this week I packaged perch into our pipeline. This required creating a program which handles the different object recognition threads in ROS. The master creates a ROS node and service which can be called from the state controllers. The ROS service, when called, organizes the input data (X and Y bounds, shelf plane height, input model STLs. Input point cloud) and then passes this to all object recognition instantiations. The algorithm outputs the model's pose on the shelf. Figure 2, below, shows the filtered bin image (left), a heat map of the depth cloud (center), and the object recognition output (right), predicting the item's position on the shelf.

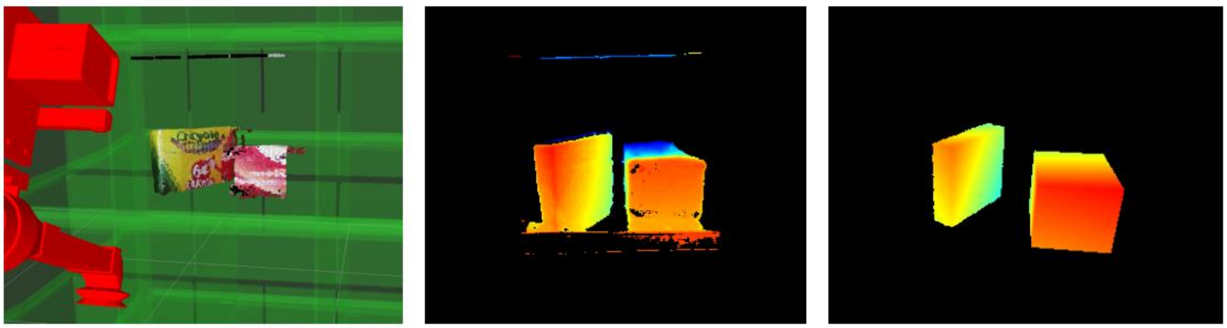

Figure 2: Filtered point cloud (left), input depth map (center), and perception results (right)

Finally, I worked on another perception approach which will perform item identification using an AlexNet architecture. Last progress review, I described this approach at a high level. This week, I worked with Alex to capture 900 images of five different items on the turntable. Some example images are shown below in figure 3. I used these images to train a neural network for item identification. The network performs identification with almost 99% accuracy on our test dataset, although there is probably some overfitting occurring. This week I will test out identification in our pipeline.

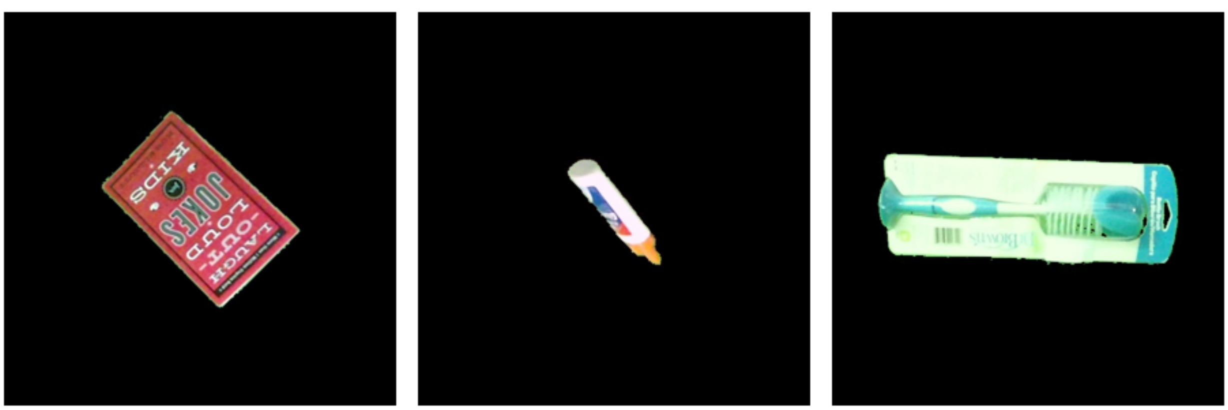

Figure 3: Training images for AlexNet item identification

On a programmatic level, the items and the competition shelf arrived. I worked with Alex to port our existing code to run on the new shelf. This included generating a model of the Kiva shelf. We also investigated the arm configuration space, and boosted the robot up 6" to have better manipulability in the top shelves.

Challenges

The biggest technical challenge over the last two weeks has been getting started with the AlexNet architecture. Caffe has had a slight learning curve, and it has been difficult to ensure that our training data is robust enough to work in our test environment. It is hard to find a good way to generate training data that is not too labor intensive. Otherwise, the other biggest challenge is just finding time to pursue both approaches, the CNN approach and geometry based approach, at once.

Teamwork

Over the last two weeks, Alex has worked extensively on adding functionality such as trajectory playback and making the arm motion control be more robust through several different planners. In addition, he did a lot of work getting the turntable setup and also transitioning to the new shelf. Feroze and Alex both focused on grasping: specifically, Feroze implemented a grasp planner that operates based on point cloud information. Alex redesigned the suction pressure sensor for improved grasp feedback. Abhishek worked on executive scripting and state control; he improved the robot's high level planning. In addition, he has worked on perception tasks with me, including pre and post processing images for training and testing. Finally, Lekha has researched rendering based methods for object identification.

Plans

By next week, the overall goal is to have multiple bin item acquisition integrated into our system. In order to accomplish this, I need to iron out a few more details. Specifically, I need to take the vision pipeline output and return a point cloud to the grasp planner. In addition, plan to have concrete numbers in regards to the learning based perception system.