

Carnegie Mellon University

The Robotics Institute, School of Computer Science

MRSD Individual Lab Report

ILR03



Team D: GetAGrip.AI

Teammates: Solomon Fenton, Jiyoon Park, Shri Ishwaryaa S V, Sridevi Kaza, Alec Trela

Author: Shri Ishwaryaa S V

Submission Date: Mar. 2, 2023

1 Individual Progress:

1.1 MRSD Project:

As I am part of the perception subsystem, these are the tasks I have completed between 2/16/2023 and 3/2/2023:

- 1) Set up testbed environment in the high bay & tested the pepper detection algorithm
- 2) Created peduncle dataset with 100 images for training U-Net architecture
- 3) Training & tuning U-Net with different backbones for peduncle segmentation

One thing I looked at but do not have results to show is:

1) Applying traditional computer vision techniques for peduncle detection

During the previous progress review, I mentioned that the YOLOv8 model that was trained for pepper detection had a saturating validation accuracy. It turned out that the model was reaching its state-of-the-art accuracy values which caused the saturation. Hence, it was safe to move on to the next task without wondering if the pepper detection algorithm could be improved.

We set up our testbed environment in the high bay by placing fake peppers on a real plant as shown in Fig. 1. The short-range and long-range testing of our pepper detection algorithm was conducted by taking videos of the pepper plant from different distances. These videos are then passed to the YOLOv8 model for inference. One observation from the tests is that the model performs better when it is closer to the peppers. In addition, it gives some false positives with a low confidence value. Therefore, we must filter the inference results to consider only the detections with high confidence values. We did not create a separate training dataset of our testbed environment because the current model which was trained on a farm environment works well here too.

A dataset of 100 images with visible peduncles was created by finding them on image search engines. VGG Image Annotator, an online image segmentation tool, was used to obtain a JSON file containing the segments. This JSON file is converted to a black-and-white mask of peduncles in an image with a Python script. This dataset is currently being used for training.



Figure 1: Testbed - a fake plant with real peppers

The neural network that I trained is U-Net. It has an encoding part that encodes an input and a decoding part that provides an output. We can switch the encoding part of a U-Net with different pre-trained networks (called the backbone) using a Python library called segmentation models. I started with Inceptionv3 as a backbone but the model was overfitting to the small dataset that we currently have. The meanIOU was only 50% on the test dataset. This is because Inceptionv3 has a large number of weights. Hence, I tried ResNet34 as a backbone because ResNets are capable of learning a shallow model from a deep network. ResNet34 also has half the number of weights. The accuracy of the peduncle segmentation went up to 75%. However, it can be seen that the result shown in Fig. 2 needs improvement. I am currently training U-Net with the ResNet18 backbone.

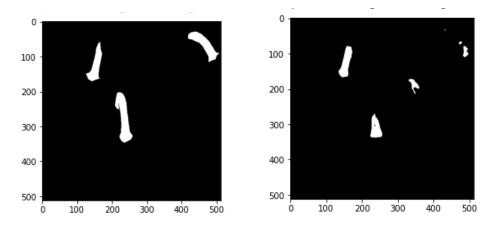


Figure 2: Ground truth mask and predicted mask of peduncle

2 Challenges:

2.1 MRSD Project:

One problem that we had is that our pepper detection algorithm had false negatives when the camera is far away from the plant. We will work on improving it but this is not necessary as long as the algorithm detects the peppers when the camera is near the plant. Another challenge that we faced is the lack of a peduncle dataset. This was overcome by manually creating one. A current issue that I am tackling is improving the performance of the peduncle segmentation model. To increase the accuracy, I will be trying different backbones and tuning the hyperparameters of the model.

3 Teamwork:

3.1 MRSD Project:

Alec Trela

Alec was able to move motors using the Dynamixel. He also created a CAD model of the cutting mechanism which was 3D printed and integrated with the existing gripper. One other task that he completed was setting us rosserial in the Arduino so that it can communicate with the ROScore running on the laptop through a serial port.

Solomon Fenton

Some tasks that Solomon finished in these two weeks are designing a spacer plate and mounting rails for the Dynamixel motors, designing and 3D printing multiple TPU gripper pads, and integrating the gripper with the cutter.

Sridevi Kaza

As I mentioned in my last individual lab report, we no longer have a xArm that works. Instead, we are planning to do our SVD with the Sawyer robot. Sri set up the robot in RViz and was able to generate trajectories with obstacle avoidance. Furthermore, Jiyoon and Sri ensured the physical robot is up and running.

Jiyoon Park

Jiyoon finished training YOLOv8 to reach an mAP of 80%. She wrote the code for predicting peppers in real-time using the model. She performed the short-range and long-range testing of the pepper detection algorithm with me.

The team went to a supermarket together, purchased peppers with the correct peduncle length, and set up the testbed. Sri and Jiyoon spent around three hours connecting the Sawyer robot with its control box and learning to communicate with the robot using a laptop. We also have alternate Fridays as demo days where we demonstrate the work completed in each subsystem in more detail to all team members. We explain the technical aspects thoroughly.

4 Plans:

4.1 MRSD Project:

Tasks that I wish to achieve before the next progress review are:

1) Train a peduncle segmentation model to meet our performance requirement

Achieving the desired accuracy for peduncle segmentation can be done by increasing the size of the current dataset and by experimenting with different models.

2) Write the main code functions in the perception pipeline - aim to cover 40%

A flowchart for the pipeline has already been created. The important functions in the pipeline will be developed using this as a reference.