INDIVIDUAL LAB REPORT 9

October 25, 2019

John Macdonald

Wholesome Robotics, Team E

Teammates:
Aman Agarwal
Aaditya Saraiya
Hillel Hochsztein
Dung-Han Lee

Contents

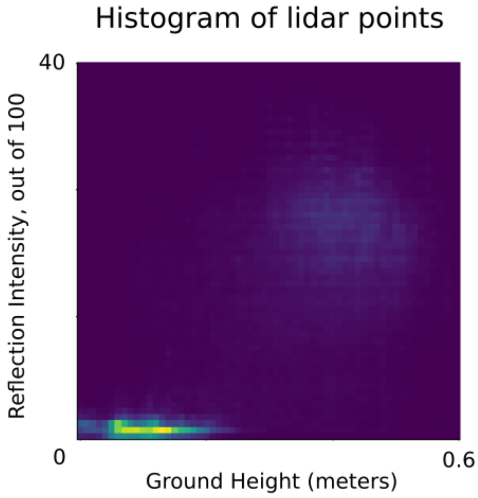

Figure 1: Histogram of lidar points in the ZI feature space, before learning.

0.1 Individual Progress

I implemented several improvements to the row perception pipeline. First, since the data from the previous field test revealed that the intensity of the points is very discriminative in terms of separating plants from ground (in this case, the weed suppression net), I implemented a weed-ground classifier using the height above the ground and the intensity of the lidar point as features. The height above the ground is obtained via the pipeline which transforms the pointcloud into the stabilized ground frame using the IMU. The point intensity is read directly from the cloud provided by the Velodyne driver. These features will be referred to as ZI point features.

Initial visualization of a histogram of the ZI features showed two very obvious modes in the distribution with Gaussian shape, which suggested that a Gaussian mixture model would be appropriate. This is shown in Figure **??**. It also suggested that I might be able to obtain good results, without have to set thresholds, or label data, via unsupervised learning on a dataset containing only these modes. I therefore implemented this solution.

The resulting clustering is shown in Figure 2.

Previously, since XYZ based thresholding was used, which was not sufficently discriminative to remove noise, RANSAC was used to fit points to the plant points. However, this lead to poor row estimates, as the rows do not actually lie on a line (they lie in a rectangular pattern), and various workarounds to address this have failure modes in which a diagonal line across the points ends up being the best "fit." With this new method, which results in less noise, a linear least squares line fit should be possible. This is equivalent to finding the major axis of the point distribution, which should go straight down the row.

The results of the new row fitting technique is shown in Figure 3.

0.2 Results

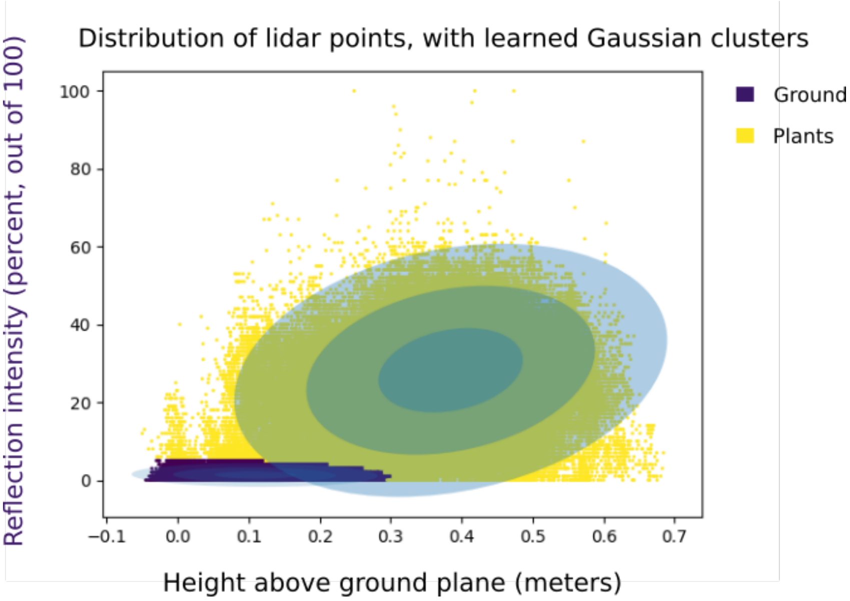

Figure 2: Results of unsupervised learning with a Gaussian Mixture model to learn the decision boundary between plants and ground. The dataset is extracted only from pointclouds inside the row, which contains only two modes in the distribution, which can be easily learned.

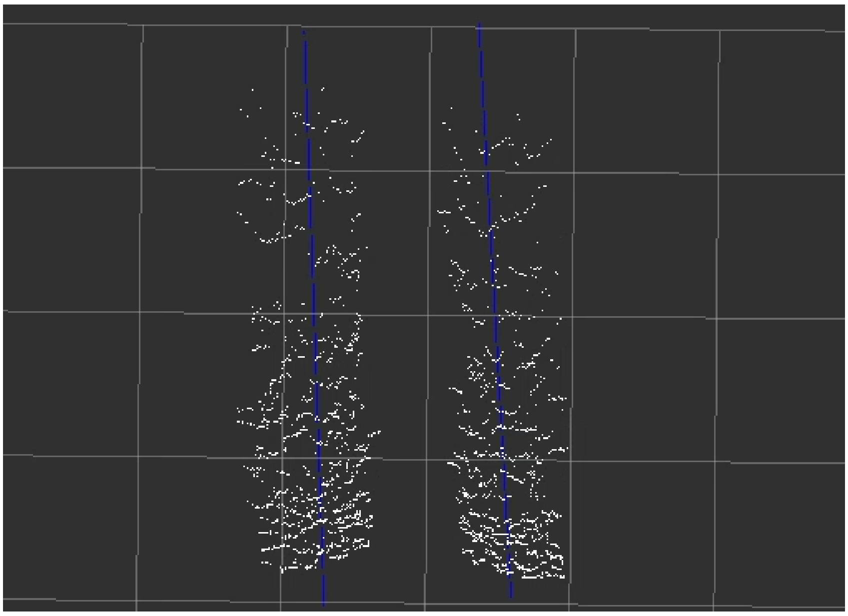

Figure 3: Rviz visualization of the new row detector. Plant classified points are shown in white and row lines are shown in blue.

0.3 Challenges

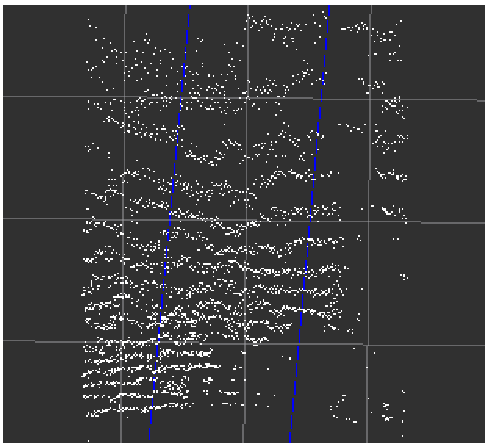

Figure 4: The GMM is not trained on data showing the grass outside the row, which has a infrared reflectively similar to the crops, so grass incorrectly classified as plants.

Since the GMM is not trained on data showing the grass outside the row, which has a infrared reflectively similar to the crops, so grass is incorrectly classified as plants, as shown in Figure 4. Currently, an additional height-above ground threshold is employed, however, it is set very conservatively to avoid interfering with the classifier output too much.

Additionally, currently there is simple cropping of the pointcloud in the ground stabilized frame based on the known dimensions of the rows in order to separate the left, right, and neighboring rows. It is navie and honestly was only done previously due to the amount of time spent fixing the particle filter. It introduces a failure mode in which points from a neighboring row are used in the fit, and pull the row line in that direction.

Potential long-term solutions for both of these issues are explored in the next section.

0.4 Teamwork

1. Hillel: UI for visualizing plant health
2. Aaditya: UI for visualizing plant health
3. Aman: Work on using `robot_localization` for fusion of VO + IMU
4. John: Improved row detection
5. Dung-Han Lee: Plant health monitoring

I helped Aman integrate his module for UTM-RTK transformation in my code. I also discussed the use of `robot`$_{l}ocalizationwithAmanandAadityatoenablefutureeasyintegration.Ia$

0.5 Plans

The issues of mis-classified ground points should be addressed. There are several solutions to this, for example only outputting row fits which have a certain expected average divergence from the center of the row line, based on the width of the crop rows. Another solution is to learn a single threshold in height and intensity and "and" the classifications together.

The issue of points from other rows degrading the fit should also be addressed. A better approach would be to cluster the plants points based on local connectivity (e.g. does it have a neighbor within X distance?), filter the clusters based on expected properties e.g. size and then fit to all such clusters. Additionally, robust error functions such as the Huber loss, which is quadratic up to a distance ϵ away from zero, and then becomes linear, could be explored.