Individual Lab Report #9

Hillel Hochsztein

Wholesome Robotics (Team E)

Teammates: Aman Agarwal, Dung Han Lee, John MacDonald, Aaditya Saraiya

October 24, 2019

Individual Progress

In the past few weeks I have continued worked on the visualizer tool for the monitoring pipeline. A key part of our deliverables is a tool for the farmers to understand the data that we collect and generate. For that reason, I have been focusing on integrating the visualizer into a completed monitoring pipeline. Specifically, we are working towards a single command line instruction that will select the desired data, inference plant health for each stored image and display the results in the final visualizer tool. The main integration tasks then were the data saving functionalities of each subcomponent, such that each subcomponent can pull information from the previous step and pass it to the next step. Specifically for the visualizer, the tool needs to select either the recently inferenced dataset, or a dataset that it has saved, due to changes made by the user. Since our three components are all run in python, we are using pickles with a standardized format to save each step. Because of this standardization, the visualizer was modified to run even off of uninferenced data (a score of -1 in either hole or disease leads to a blue marker). This change is useful in the long term, because if no result can be determined for a given picture, the picture is flagged and handled, rather than causing an error.

I also upgraded the clustering to handle integration. The clustering is accomplished via line fitting and then matching images to the closest drawn point. In order to integrate, this algorithm needs to work for multiple lines, while maintaining the same output to the actual gui. Doing so required only slight modifications to the base algorithm (sorting the data points by row, and clustering one row at a time, before combining the results back into a globally consistent output), but as always, upgrades and integration required large amounts of debugging.

Due to these efforts, the visualizer was completely integrated into the pipeline. I inputted its launch commands into the bash script that runs the complete pipeline, and it worked completely, albeit not for the inferenced data, as we discovered a bug in the save commands for the inference pipeline. Barring this issue, we have completed intgration for the pipeline, and are just about done with our MVP.

Finally, I introduced a few relavent features, necessitated by the clustering, that will ease usablity. The first of these was to cause the navigation buttons to deactivate (and visually show this) when at the limits of the list of images. This allows the user to know that they are at the beginning or end of the list. The other major elements are an index heading (telling the user which image of the cluster they are looking at), and a status bar that shows the health of all the images in the cluster. These allow the user to more intuitively understand the changes they are making with the gui. Accomplishing these tasks were

not technically complex, however they required a fair degree of research and debugging. Most importantly, these tasks were fairly separate from the integration, and were not blocking any work the rest of the team needed to do. This was important because of my restricted schedule for the past 2 weeks.
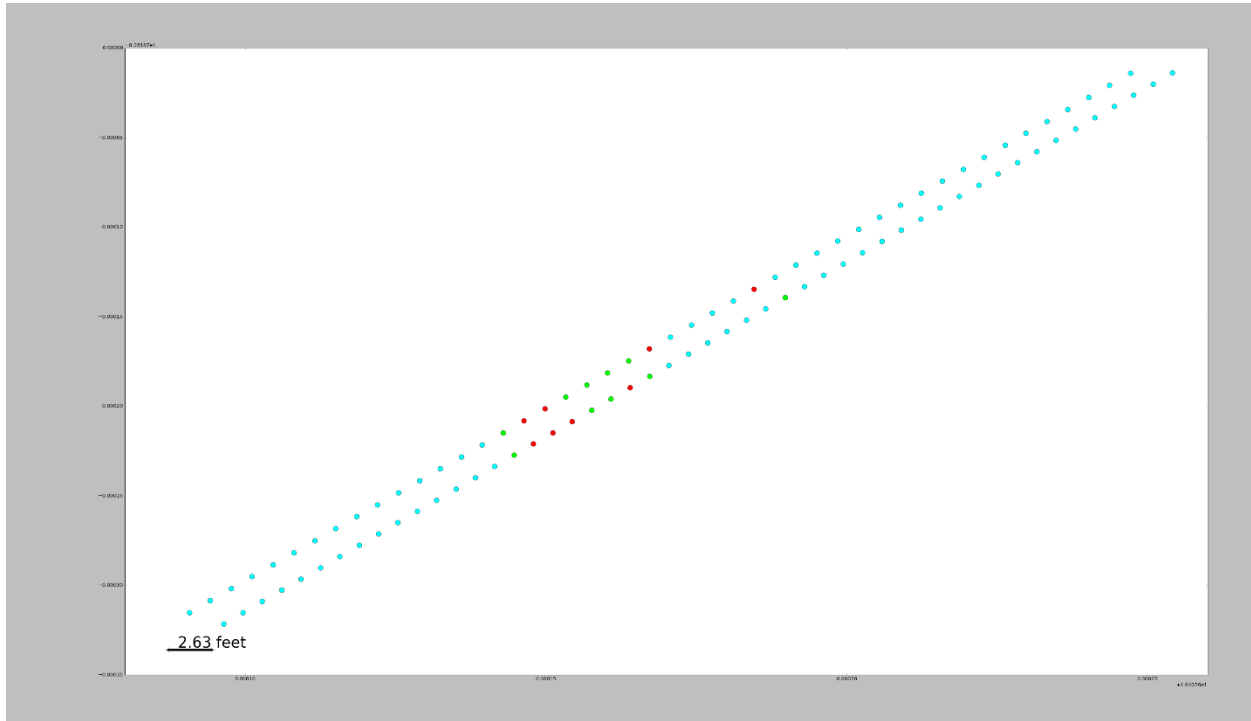

Figure 1 Clustered version of a Multiple row of plants. Green demonstrates healthy, Red unhealthy, and blue uninferenced (this was run on a test set, that was not actually run through the inferencing pipeline).
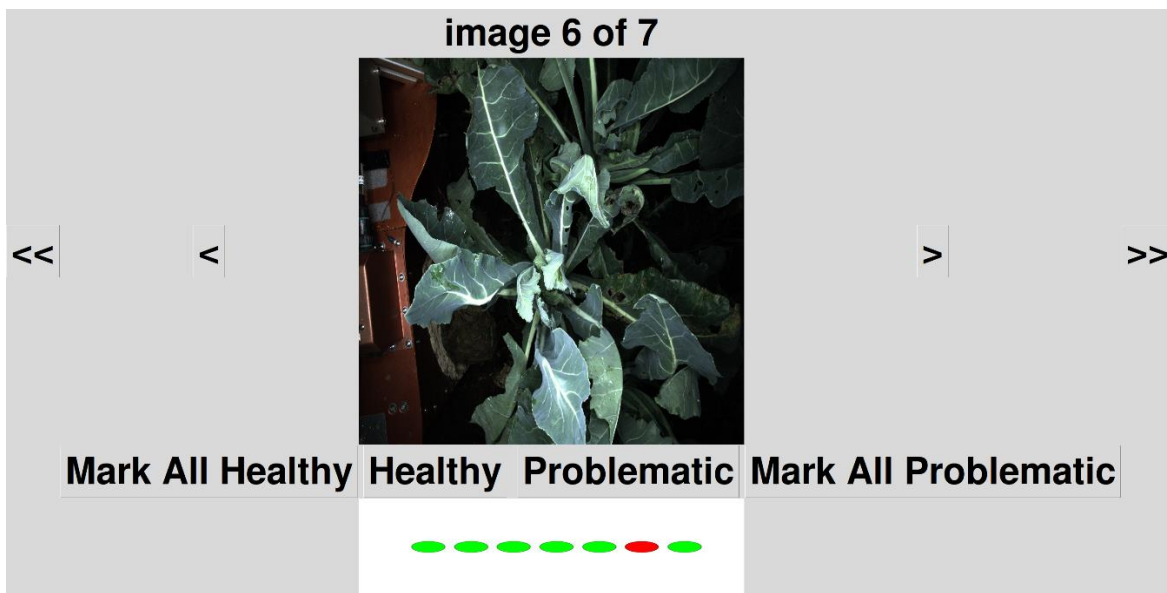

Figure 2 GUI includes readout for image number within the cluster, total cluster status bar, and navigation buttons that deactivate when unavailable (i.e. at the first or last image)

Team Progress

We have been working towards delivery of our MVPs both for visualization and for navigation. For visualization we have been working on communicating and implementing interfaces between different components as we continue to improve the various components (the classification and the visualizer). For navigation we have worked on fusing different row detection and row navigation techniques (such as the lidar based detection, RTK guidance, and visual odometry) to fix issues we found with our row navigation.

Challenges

The biggest challenge I have faced has been time management. Because of the Jewish holidays I have been unable to work or even communicate with my team for more than 7 days over the course of the past two weeks. Fortunately, the team has been very understanding about coordinating meetings to accommodate my availability, and I have taken deliberate actions to manage my remaining time.

Teamwork

Aman has worked on the row navigation techniques, and specifically sensor fusion

Aaditya has been working on visualizer integration and image exposure acceptance for the crop facing camera.

Dung Han has been working on evaluating the plant health model and fine tuning it, including comparing the current model with other architectures.

John has been working on the localization and navigation nodes.

Future Plans

Wholesome Robotics

In the coming weeks my plan is to finish integration of the visualizer MVP. This mostly involves debugging compatibility and interfaces between the different parts of the pipeline. As mentioned before, the interfacing is a mostly completed project, meaning the remaining time will be spent implementing the final farmer requested features (the comment bar and row labelling are the primary requirements, other features are considered stretch goals).

The wheel guards, which have been deprioritized until now, need to be fabricated based on the preliminary design I created a few weeks ago.