Progress Review 11

Individual Lab Report 10

Aaditya Saraiya

Nov 7, 2019

Team E:

Wholesome Robotics

Teammates:

Aaditya Saraiya

John Macdonald

Dung-Han Lee

Aman Agarwal

Hillel Hochsztein

Individual Progress

Capstone Project

Team E is creating an organic monitoring robot which has to autonomously navigate through crop rows. For this phase of the project, I had two main tasks which can listed as follows:

1) Integrating the complete monitoring pipeline and testing the functionality of the complete pipeline.

2) Working with Aman on utilizing the robot_localization package and testing initial results in a controlled environment.

Monitoring pipeline integration

Previous status

Before this PR, the status of the monitoring pipeline can be summarized as follows:

1) The ROS Bag parser which takes images and associates it with image locations was integrated with the deep-learning based plant health detection pipeline. This pipeline is based off the Mask-RCNN neural network.

2) The visualizer was integrated with the full pipeline for single ROS Bags. However, the clustering for multiple ROS Bags had been independently tested, but not integrated with the complete pipeline.

3) As discussed in previous ILRs, an exposure checker script was developed to check if more than 75% of the images from a collected ROS Bag are usable images or not. Usable images are currently defined on basis of visual inspection. The classification of images as usable or not usable was done by analyzing the image intensity histograms.

Progress updates

The image exposure script was integrated with the complete pipeline and tested. For the three tests bags picked from data collected from the previous field visits, the image exposure script was passed successfully. However, there were some integration issues faced, which have been discussed in details in the Challenges section.

My other tasks in this effort involved facilitating the integration of the visualizer with rest of the pipeline. Primarily, I was collaborating with Hillel for this task. During the integration process, we faced certain issues related to .pickle file formats between Python versions, which have been described in the challenges section.

Overall, the conclusion on this part was that the demo was not successful up to the targets set for this pipeline. For the monitoring pipeline, it has been observed previously that integration takes considerable time due to multiple members involved in the development process. These issues will be corrected for the upcoming PR.

Sensor Fusion for Localization

State of pipeline before PR

Before this progress review, Aman and I had done some initial work setting up the robot_localization package. An Unscented Kalman Filter based pipeline was utilized to provide local fusion estimates at higher rate by fusing in visual odometry and IMU data. An Extended Kalman Filter was utilized in order to provide global fusion updates by fusing in the GPS data. The pipeline was tested on ROS Bag files and it seemed to

give decent results for translation along the row. However, the filter used to diverge as soon as the robot turns on the field. The conclusion from these tests were that the filter is not able to handle changes in the yaw estimates.

Results with first field test

Unfortunately, the development time for this module conflicted with the field visit test time with the field visit falling in too early. We tested the UKF+ EKF based sensor fusion module during the field visit. However, it was observed that the filter was diverging from the start and was not performing well even for translation along the row. Hence, the initial conclusions were deemed incorrect.

Recognized issues

In order to debug the issues from the first phase of work with this module, the following steps were listed down.

1) Checking and confirming that the IMU was oriented using the ENU convention and not the NED convention.
2) Check and confirm that the ZED pose estimates were in the correct coordinate frame.
3) Write a covariance estimator node which collects data from the IMU and ZED and finds the initial covariance for the sensors.
4) Perform extensive indoor testing on the UKF node to confirm correct functioning. Make sure to do all the debugging directly on hardware.

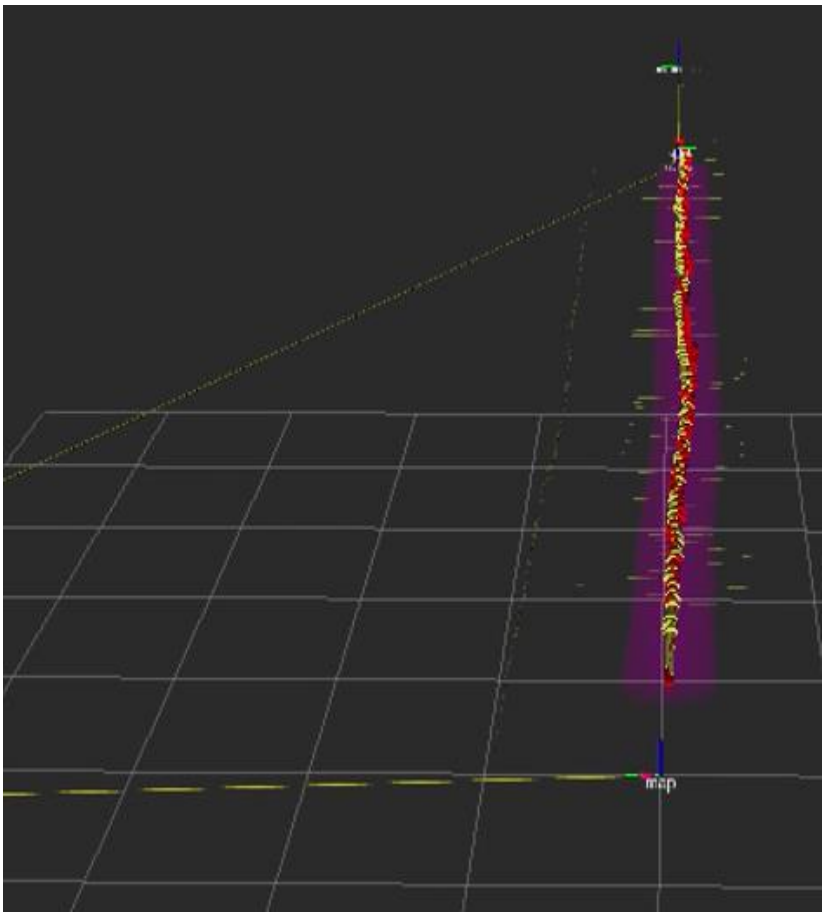

Figure 1: Divergence in the UKF + EKF based node on ROS Bag data

Added improvements and testing hypothesis

Specifically, I worked on writing the script to estimate initial covariance from the IMU and ZED sensor. This script takes a parameter as input which decides the minimum number of readings which need to be taken from the IMU and ZED before estimating the covariance from the data. The robot was kept idle for a few seconds and the co-variance data was obtained. The cross covariance terms from the data were ignored for simplicity.

On further testing, it was confirmed that the IMU and ZED pose and orientation data was according to the convention required by the robot_localization packages. One of the major issues which was recognized from the testing process is that the ZED sensor readings were not time-synced with the other sensor readings from the robot.

Initial coarse tuning of the EKF and UKF parameters with using the initial covariance estimated from the script, it was observed that the localization estimate has converged for indoor environments. Initial tests show error of ~10 cm for 5 meters of trajectory. Further experiments in the coming PR will focus on documenting the effect of tuning parameters on the pose estimation error.

Interacting with data labelling companies

Previously, all the labelled data utilized by the deep learning pipeline had been manually labelled by Dung-Han Lee. However, for the new data, keeping the time constraints in mind, the team required us to outsource the labelling process for the deep learning pipeline. I was in charge of contacting the companies for the data labelling. After some initial options, we decided to choose Scale.ai. We are currently in contact with the company representatives and the data will soon be sent out for labelling.

Challenges

This particular progress review involved a lot of challenges, for both the sensor fusion pipeline as well as the integration task. The challenges section has been divided into two parts to elucidate the challenges faced with both the pipelines.

Sensor Fusion Pipeline

For the sensor fusion pipeline, the major challenges were involved with correctly using the robot_localization package. Our efforts to get initial results from the package before the field visit were not successful. However, after the field visit, we developed a structured approach to debugging the package and we were able to get successful localization for indoor environments (after the presentation on Wednesday). A lot of initial time working with the package was wasted in tuning parameters, when the actual issue revolved around communication issues between the master and slave computers. Hence, testing the package directly on hardware would have exposed the issues earlier as compared to testing on ROS Bags.

Monitoring pipeline integration

For the monitoring pipeline, the goal for this progress review was to perform a complete integration of the monitoring pipeline. However, there were some specific problems which prevented us from achieving this goal. Firstly, there were compatibility issues with the reading the pickle file between Python 2 and Python 3. This lead to some features in the visualizer to be disabled. Secondly, there was some poor version control practices on my behalf which led to the exposure checker file to be deleted. Hence, we were not able to showcase this during the PR. This issues reiterates the importance of version control and the need for regularly pushing code onto Github.

Teamwork

John Macdonald

John worked on creating a new version of the row detector and integrating it with the particle filter which was developed as a part of last semester. He also did some initial work on creating a deep learning based pipeline to perform row segmentation with LIDAR data.

Aman Agarwal

Aman worked primarily on the robot_localization package to perform sensor fusion with the IMU, visual odometry from the ZED camera and GPS data. He worked on setting up the UKF and EKF launch files for local and global sensor fusion respectively.

Hillel Hochsztein

Hillel worked on adding additional features for the plant health visualizer. Specifically, he developed a note-taking feature which allows farmers to add specific text notes from visual inspection of the visualizer output.

Dung-Han Lee

Dung-Han Lee performed further analysis and tests with the UNet based model. He validated the UNet model based on the old data. The UNet model achieved the required precision and recall for the old data.

Future Plans

Team

With the monitoring as well as navigation parts of the pipeline coming to maturity, this next PR will focus on solving the existing problems in the pipeline and showcase initial MVPs for both monitoring and navigation pipelines. The future team goals for the upcoming weeks can be summarized as follows:

1) Fixing all the pending issues for the final integration for the monitoring pipeline.
2) Adding an option to run the new UNet based model with the monitoring pipeline for plant health detection.
3) Demonstrating the initial set of successful results for localization and autonomous navigation through the crop rows at Rivendale farms.
4) Improving the localization results by fusing in LIDAR row detection data.

Individual

The future individual tasks which are planned for the coming weeks can be summarized as follows:

1) Developing a sensor model to fuse in LIDAR detection data with the robot_localization package.
2) Tuning the EKF and UKF parameters to reduce drift with the sensor fusion pipeline.
3) Facilitating the final integration process for the monitoring pipeline.