ILR05 - Progress Review 4

Justin Morris
Teammates: Awadhut Thube, Alex Withers
Team G: The Pit Crew

September 17, 2020

1 Individual Progress

Due to COVID-19, as well as the normal process of re-acclimating to the new semester, we set achievable goals for our first progress review of the fall. Conveniently, Alex secured an internship position working for Red Whittaker on the Pit Ranger project, which allowed him to advance the project significantly over the summer. Because of this, one of our two goals for Progress Review 7 had already been completed by the time I returned to Pittsburgh in late August.

The goal which required additional work was to develop a more sophisticated heuristic for determining the distance from the rover to the pit edge. In previous testing, we had simply counted the number of points in the point cloud generated from our stereo image data, and stopped the rover when this number dropped below some threshold. We recognized that this heuristic would not be sufficient for the rover to operate safely on the moon. Using our domain-specific knowledge and the data we had collected in the spring, we needed to come up with a more robust method of detecting edges in the rover's perception space.

I came up with the idea that we are currently planning to implement. I knew that in many video games, the virtual environments were simulated as polygons, triangular planar surfaces that fit together to form approximations of 3D shapes. I reasoned that if we could convert our point cloud into polygons, we could make judgments about the slope of each planar surface and use that to identify navigable terrain. In order to accurately model the terrain, the polygons would have to be generated by clustering points in the cloud according to both their location in space and their planar characteristics. This would ensure that points near a sharp edge in the environment would be clustered into separate polygons on each side of the edge. Then, if the normal unit vector of a given polygon had a Z component that was below some threshold, that would indicate that the slope in that area was too steep for the rover to traverse. This would give the robot the ability to identify and avoid dangerous slopes.

This concept, however, did not provide a solution for identifying sheer brinks. We knew from our previous tests that as the robot approaches a drop-off in the terrain, that area will be perceived as an empty space in the point cloud. Therefore, I theorized that we could compare our polygonal representation of the area perceived by the rover's cameras to the area of the terrain that we would expect to perceive on flat ground. If our polygonal representation contained empty spaces where we would expect to see points, we would know that that section of the environment was a pit or other brink. Beyond the area that the robot can perceive, it would obviously be impossible to determine whether a pit is present, but as long as the rover updates its point cloud rapidly enough in comparison to its movement speed we can catch brinks as soon as they appear in the perceivable area.

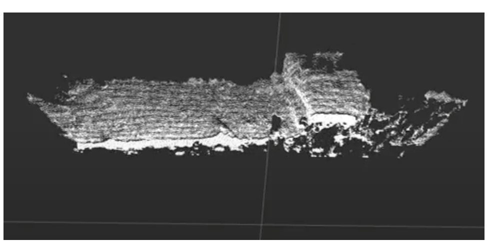

(a) Generate point cloud from stereo data.

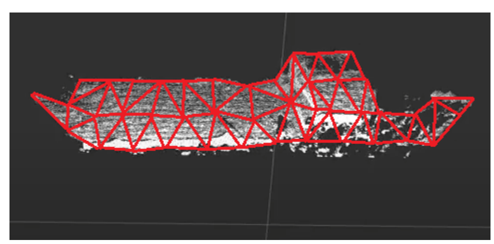

(b) Cluster points into polygons.

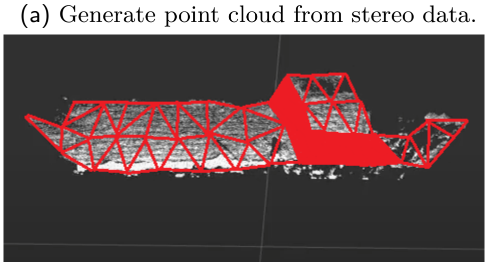

(c) Identify polygons with untraversable slopes.

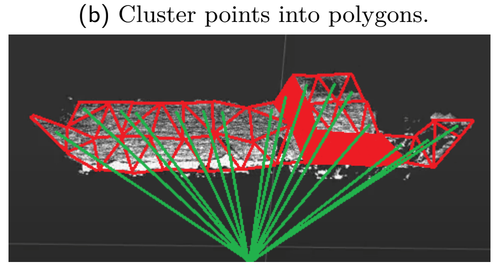

(d) Measure distance to the far edge of the polygon lattice.

2 Challenges

Having an idea for how to approach this problem is the first step, but implementing a working version of the proposed heuristic is a much more daunting task. I anticipate that development of this heuristic will occupy a significant portion of the semester, and will hopefully line up well with our team's proposed goals of being able to identify one new type of terrain at each of the next four progress reviews. A particularly challenging aspect of the heuristic will be the polygon generation. It is very important that the polygonized terrain accurately models the actual shapes of the surface, so that the rover's will not reach erroneous conclusions about how to navigate its environment. Finding a method for converting a point cloud into accurate polygons will be difficult. Although many disciplines use polygon representations of environments, I do not know what prior work has been done on generating those representations using a point cloud.

This kind of technical challenge is, in a sense, a welcome break from the more administrative challenges that we have encountered over the past several months. COVID-19 continues to be a major challenge to the project, making it more difficult for our team to access on-campus resources, work together in close proximity, and perform tests with our surrogate rover in field environments. Luckily, Alex has made significant progress on the simulator over his summer internship, and his expertise will be invaluable in allowing us to make progress without the need for physical hardware. However, a larger emphasis on simulation provides a challenge of its own, as Awadhut and I must implement the simulator on our own computer systems and become as familiar with it as Alex has.

3 Teamwork

As previously stated, Alex's work on the simulator over summer has been hugely beneficial to the project as a whole. Since we returned for the fall semester, he has assisted in bringing me and Awadhut up to speed on use and development of the simulator. Awadhut has successfully run the simulator on his system, although I am still working on setting the simulator up to work properly for myself. One useful feature that Alex has developed is a smaller simulation environment which contains examples of the four different slope categories that we intend to recognize. By simulating the rover in this environment, we can validate the terrain recognition functions as we develop them.

Alex has also developed connections with other people working on the Moon Ranger and Pit Ranger projects through his summer internships. These connections will be crucial in allowing us to integrate our work with the other subsystems and developments that make up Pit Ranger. We will maintain those relationships so that we can work in concert with them throughout the upcoming semester.

The entire team collaborated on the development of the pit distance measurement heuristic. Although the base concept was mine, I refined and developed it with the assistance of Alex and Awadhut, and they each used their specific knowledge of the simulator and the perception subsystem to suggest alterations and confirm the feasibility of the idea.

4 Plans

Alex will continue his work from over the summer on the planning subsystem. He will refine his existing code and verify it through testing in the simulator. He will also integrate the localization and mapping code developed by other members of the Moon Ranger team so that we can gain the functionality that they have developed. Meanwhile, Awadhut and I will work on developing the terrain recognition and pit distance measurement heuristic. Awadhut, who has previously been the expert on the perception subsystem, will focus on converting the point cloud into polygons. I will create the code to assess the traversability of each polygon, and determine if brinks are present in the rover's field of view and where the brinks are in relation to the rover.

Initially, all of our development will be done in the WeBots simulation. We will plan to test on a physical rover surrogate later in the semester, once we feel confident in the robustness of our code. The physical rover tests will be smaller in scope than the set of functions that we will demonstrate in the simulator at the end of the semester, but they will act as a proof of concept that our system can function in a true field environment.