

Experimental Results in Using Aerial LADAR Data for Mobile Robot Navigation

Nicolas Vandapel Raghavendra Donamukkala Martial Hebert

The Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213

vandapel@ri.cmu.edu

raghu@ri.cmu.edu

hebert@ri.cmu.edu

Abstract

In this paper, we investigate the use of high resolution aerial LADAR data for autonomous mobile robot navigation in natural environments. The use of prior maps from aerial LADAR (LAser Detection And Ranging) survey is considered for enhancing system performance in two areas. First, the prior maps are used for registration with the data from the robot in order to compute accurate localization in the map. Second, the prior maps are used for computing detailed traversability maps that are used for planning over long distances. Our objective is to assess the key issues in using such data and to report on a first batch of experiments in combining high-resolution aerial data and on-board sensing

1 Introduction

Autonomous mobility in unstructured, natural environments is still a daunting challenge due to the difficulty in analyzing the sensor data from mobility sensors, such as stereo or laser rangefinders. Recent developments make it possible and economical to acquire high-resolution aerial data of an area prior a ground robot traverses it. Although the resolution of conventional Digital Elevation Maps (DEMs) is too limited to be used effectively for local robot navigation, high-resolution aerial mapping opens the door to pre-processing a terrain model at a resolution level comparable to that of the on-board sensors.

In this paper, we investigate the use of high resolution aerial LADAR data for autonomous mobile robot navigation in natural environments. The use of prior maps from aerial LADAR survey is considered for enhancing system performance in two areas. First, the prior maps are used for registration with the data from the robot in order to compute accurate localization in the map. Second, the prior maps are used for computing detailed traversability maps that are used for planning over long distances. Our objec-

tive is to assess the key issues in using such data and to report on a first batch of experiments in combining high-resolution aerial data and on-board sensing.

In the application considered here, the typical mission scenario aims at performing way-point navigation over hundreds of meters in a rough terrain cluttered with various types of vegetation. The ground vehicle system, from General Robotics Dynamic Systems (GDRS), is equipped with a three dimensionnal (3D) laser sensor, an inertial navigation unit and a Global Positioning System (GPS). Overhead data - including LADAR data and multi-spectral imagery (not used here) - is provided prior to the mission.

Using aerial LADAR data poses two main challenges: the volume of data and the nature of the terrain. The aerial LADAR data contains typically 44 millions of 3D points each associated to an identifier and a reflectance value and covers $2.5 \text{ km} \times 3.0 \text{ km}$, which is an enormous amount of data to use effectively on-board a robot. In addition, the data includes two components of the terrain - the vegetation cover and the terrain surface, which need to be discriminated from one another. The work reported here is a first step toward using this type of high-resolution aerial data in conjunction with data from a robot's mobility sensor.

The paper is organized as follows: The second section presents the type of data we had access to and the sensors we used to collect them. The third section deals with the details of vegetation filtering, using ground and aerial LADAR data. In Sections 4 and 5 our work on vegetation filtering is applied to two problems: robot position estimation using 3D terrain registration and path planning.

2 Data Sets

This section introduces the various data sets used to produce the results presented in this article as well as the sensors that collected them.

2.1 Data collection

The data used for this work was collected on three different sites characteristic of very different types of terrain: the AP Hill site is a wooded area with large concentrations of tree canopies; the Yuma site is a desert area with scattered bush; the Pickel Meadow site is an alpine area, which includes a substantial amount of vegetation and, more importantly, large variations in terrain slope. The data used in this article was collected during these three field tests; unless explicitly noted, these results were obtained after the field tests. Prior map data was provided to us in for the AP Hill and Yuma data sets.

2.2 Sensors

We used the GDRS mobility laser scanner (similar to the sensor described in [16]), mounted on a ground mobile robot, and a Saab TopEye mapping system mounted on a manned helicopter [1] to collect data.

Ground sensor. The ground LADAR is mounted on a turret sitting at the front of the ground vehicle. The vehicle is built using the chassis of a 4 wheel drive All Terrain Vehicle (ATV). The laser has an effective range of 20 meters and a 7.5 cm range resolution. In the configuration used in these experiments, its field of view is $90^\circ \times 15^\circ$. In addition, the sensor can pan and tilt by $\pm 90^\circ \times \pm 15^\circ$ for increased terrain coverage.

Airborne sensor. The aerial mapping system is operated at 400 meter above the ground. The laser beam is deflected by an oscillating mirror which produces a Z-shaped laser track along the flight path. The range resolution is 1 cm and the point position accuracy varies between 10 and 30 cm, depending on the altitude. The laser records two echoes per pulse (first and last), but only objects taller than 1.8 meter will produce two echoes. For each demonstration, the helicopter flew over the test area along several directions to produce higher point density, 4 to 10 points per square meter.

3 Vegetation filtering

In this section we justify the need for removing the vegetation from the raw data, we review the issues of perceiving vegetation with laser sensors, we present a state of the art of the techniques commonly used, and we explain the two methods implemented for the aerial and the ground LADAR data.

3.1 Motivation

Vegetation constitutes a major challenge for robot navigation for the following reasons: 1) Vegetated areas are unstable features, susceptible to natural changes and human activities, 2) It is difficult to recover and model the shape

of trees and bushes, 3) Vegetation can prevent the detection of hazards such as trenches or rocks, 4) Vegetation such as trees and bushes might constitute a hazard, whereas grass might not and overhanging branches are sometimes difficult to detect, 5) Vegetation prevents the detection of the terrain surface, which is used in terrain-based localization, and is required when the canopies produce GPS drop-off.

3.2 Vegetation and active range sensors

Issues related to the interaction of an aerial LADAR sensor with ground cover have been well studied in related work and are summarized here.

Sensing. The foliage penetration rate will depend on: the canopy density (winter preferable); the scanning angle (nadir preferable); the laser footprint on the ground (small aperture, low altitude preferable). The signal returned might contain several echoes, one for the canopy and one for the ground.

Data interpretation. A change in elevation can be interpreted as a change in the ground cover (terrain vs. trees) or as a change in the terrain elevation. Because of the range dispersion when the laser beam hits multiple objects at different depths, the range measurements will be erroneous, depending on the flight path and the scanning angle.

Ground scans versus aerial scans. Ground scans are less sensitive to the effects described above because the laser range is shorter (dozens vs. hundreds of meters) and so the beam footprint on the target is smaller. Another difference is the geometry of the scene. Aerial LADAR will most likely map the ground and the top of the canopy. Trunks and small obstacles will not be perceived. With a ground LADAR, the bottom of the canopy, small obstacles and trunks will be seen. Trunks will produce range shadows, occluding a large part of the terrain. The last difference will be the density of points, with a difference of two orders of magnitude (10 versus 1000 per square meter).

3.3 State of the art

Filtering LADAR data has been mainly studied in the remote sensing community with three objectives: producing surface terrain models [11] (in urban or natural environment), studying forest biomass [12], and inventoring forest resources [6]. To filter LADAR data authors used linear prediction [11], mathematical morphology (grey opening) [4], dual rank filter [13], texture [5], and adaptive window filtering [15]. All these methods are sensitive to the terrain slope. In the computer vision community Mumford [8] pioneered the work for ground range images. Macedo [14] focused on obstacle detection among grass.

3.4 Methods implemented

We implemented two methods to filter the vegetation. The first one takes advantage of the aerial LADAR capability to detect multiple echoes returned per laser pulse emitted. The ground LADAR does not have this capability, so a second filtering method had to be implemented.

Multi-echoes based filtering. The LADAR scans from multiple flights are gathered and the terrain is divided into $1\text{ m} \times 1\text{ m}$ cells. Each point falling within a given cell is classified as ground or vegetation by k-mean clustering on the elevation. Laser pulses with multiple echoes (first and last) are used to seed the two clusters (vegetation and ground respectively). Single echo pulses are assigned initially to the ground cluster. After convergence, if the difference between the mean value of the two clusters is less than a threshold, both clusters are merged into the ground cluster. The clustering is performed in groups of 5×5 cells centered at every cell in the grid. As we sweep the space, each point is classified 25 times and a majority vote defines the cluster to which the point is assigned. This method has been used to produce the results presented in section 5.

Cone based filtering We had to implement a new method for filtering the vegetation from ground data. Our approach is inspired by [17] and is based on a simple fact: the volume below a ground point will be free of any LADAR return. For each LADAR point, we estimate the density of data points falling into a cone oriented downward and centered at the point of interest, as shown in figure 1.




Figure 1: Cone based filtering. Point: point to be clustered as ground or vegetation. θ opening of the cone. ρ minimum elevation before taking into account points

While the robot traverses a part of the terrain, LADAR frames are registered using the Inertial Navigation System (INS). The INS is sensitive to shocks (e.g., a wheel hitting a rock), which causes misalignment of consecutive scans. In order to deal with slightly mis-aligned frames, we introduce a blind area defined by the parameter ρ (typically 15 cm). The opening of the cone (typically $10\text{-}20^\circ$) depends on the expected maximum slope in the terrain and the distribution of the points. Figure 2 illustrates an example of points distribution for a tree as seen by an aerial and a ground sensor. In the latter case, the trunk will occlude

the terrain surface on top of which the LADAR can perceive the tree canopy. If the cone opening is too narrow, the vegetation point will be mis-classified. This case also illustrates the fact that by nearly taking the minimum elevation in each cell of a gridded terrain, we cannot recover the ground surface.




Figure 2: Points distribution in a tree for an aerial LADAR sensor (left) and a ground sensor (right)

This approach has been used to produce the results presented section 4. Our current implementation filters 67,000 points spread over $100\text{ m} \times 100\text{ m}$, in 25 s on a Pentium III, 1.2 GHz.

4 Terrain registration

In this section we present briefly our method for 3D terrain registration and we show registration results. These results are obtained in an area in which registration would not be possible without the vegetation filtering and ground recovery algorithms described above.

4.1 Terrain registration method

The objective of terrain registration is to recover the vehicle position in the map by matching a local map from 3-D data from on-board sensors with 3-D data from the prior map. The core of our registration approach involves the computation of pose-invariant surface signatures in the neighborhood of feature points in both the robot and the prior map. Correspondences between the features are established by comparing the signatures, and the most consistent set of correspondences is retained for computing the registration transformation. An initial description of this class of approaches can be found in [10]. Initial extensions to terrain matching are described in [7]. Details of the current version of the map registration approaches are described in [18]. Importantly, this approach does not require accurate prior knowledge of vehicle pose. In fact, we

have performed registration with 20m initial error in robot position and $\pm 10^\circ$ error in orientation.

Key to the current approach is the automatic selection of interest points or feature points in the terrain maps. This is challenging because the points must be selected in a consistent manner between the aerial and ground data. The approach can be summarized as follows. Points are selected using three criteria computed from the configuration of the 3-D surface in the neighborhood of each point. The first criterion uses the configuration of range shadows in the vicinity of the point; essentially, points with low density of range shadows in their neighborhood are selected in priority. The second criterion evaluates the amount of variation of the terrain surface in the vicinity of each candidate point, so that only points with sufficiently curvy terrain are selected. Finally, the signatures are analyzed so that only those points whose signatures contain enough information are retained. Given a 3-D map constructed from the on-board sensors, the algorithm that combines the three criteria extract a small set of points that are used as the feature points for matching. Andrew Johnson introduced a more elaborated but more expensive landmark point selection strategy in [9]. By contrast, our implementation runs on-board the mobile robot.

We present the performance evaluation below for the registration using the data described above. The key conclusion of the experiments is that vegetation filtering and ground recovery is critical to reliable terrain registration for two reasons. First, the vegetation generates very different data sets when viewed from the vantage point of the robot as opposed to the vantage point of the aerial sensor. This is true in obvious cases such as tree canopies, but this is also true of more mundane structures such as bushes. Second, this problem is compounded by the fact that selecting features for registration is based on the amount of variation in 3-D data around each point, which causes feature points to be selected near vegetation areas, the least reliable for matching. For example, bushes and trees tend to be the most prominent 3D structures in the scene, occluding other terrain surface features, such as rock formation and ledge. Furthermore, the last stage in registration involves minimizing the registration error between the two 3-D data sets, after transformation by the transformation computed from correspondences between feature points. Because of the stochastic nature of the data in vegetation areas, this minimization is, in practice, unreliable.

4.2 Yuma Data

Ledge Example. In Figure 3, we present an example of air-ground registration. The area registered does not contain any vegetated areas, but the aerial data does. Without filtering the vegetation, the points retained for performing the registration were mostly located in the wash, vis-

ible in the rear of the scene. The green surface represents the aerial data, it is $100\text{m} \times 100\text{m}$. The hill that is visible is roughly 10 m high and the ledge cliff is at most 3 m high. The data was collected during the demonstration performed at Yuma, in may 2002.




Figure 3: Air-ground registration in the ledge course. In red the ground data. In green the aerial data, the vegetation has not been filtered for the rendering but has been for the points selection

Wash. In this second example the ground data contains bushes and trees. Figure 5-(a) presents an aerial view of the scene with the robot trajectory overlaid, the north is oriented up, and the east toward the right. The robot drove 9.20 m westward. Trees and bushes are darker than the terrain.

The picture 5-(b) was taken from the top of the ledge course and the wash is visible on the right side. Aerial and ground data are represented as points clouds figure 5-(c) and (d) respectively. The filtering method works correctly even in the presence of a steep slope, as shown in the side of the ledge figure 5-(c). The vegetation points are plotted in color while the ground points are in white. Table 1 contains statistics about the results of the vegetation filtering for this example.

One added difficulty in this example is that, after filtering the vegetation in the robot data, the terrain surface contains a number of empty areas with no data, termed “range shadows”. This is because in many cases there is not enough data to recover the ground surface after filtering. This makes the feature selection more difficult because we cannot use the range shadows as a criterion to reject potential feature points. In the aerial data, on the other hand, the ground surface can be recovered more reliably.

Figure 5-(e) contains the aerial data without vegetation (in grey) registered with the ground data without vegetation (in red). Table 2 summarizes statistics about the registration procedure.

4.3 Pickel Meadow Data

In the example illustrated in Figure 6 we tried to register two ground scans taken apart in a highly cluttered environ-

Raw data	# points			Veg. Filt. (s)
	Total	Veg.	Ground	
Aerial	50,000	9,583	40,417	15
Ground	541,136	387393	153,721	117

Table 1: Vegetation filtering of the raw data

Mesh	Res. (m)	Size ^{##} (m×m)	# of vertices		Time (s)
			total	kept	
Ground	0.27	40×25	1252	400	2.3
Aerial	1.2	100×100	9000	637	

Table 2: Registration of the surface meshes

ment. The bounding box of the two registered ground scans is 40m×40m. We manually counted 30 trees in this area, in which only 10 percent of points perceived by the LADAR are part of the ground. The registration was possible because of the 80 cm rocks visible at the edge of the scan at the center of figure 6-(b). Figure 6 presents a side view of the two terrain surfaces (without vegetation) registered. As we can see on the left side, the registration is slightly off in rotation (two degrees). One problem which arises in this example is that the vegetation filtering method has not been able to correctly filter the stump of some trees visible in one scan and not in the other one.

5 Path planning

In this section, we describe the use of prior maps from high-resolution aerial data for path planning. We consider the problem of planning from a start position to a distant goal point through a set of intermediate waypoints. We evaluate different options for processing the prior map data and show the effect of vegetation filtering on the performance.

5.1 Approach

Traversability map. Our task is to provide a path between way points based on aerial LADAR data. Figure 4 contains a flow chart of our approach. Starting with the raw data we first divide the space into 1 m × 1 m cells. Points are then segmented into two clusters: vegetation and ground. For each cell the ratio between the number of vegetation points and the total number of points defines a confidence criterion on the terrain reconstructed below the canopy. We called it the *vegetationness*. This new criteria allows the path planner to consider new trajectory, usually prohibited because vegetation is not treated a priori as an obstacle.

The traversability at each location in the map is evaluated using the standard approach of convolving a vehicle model with the elevation map. More precisely, using the ground points we compute height traversability maps (one every 45° in heading) as follows: 1) we interpolate the

position of each tire on top of the ground surface, 2) we fit a plane and extract the current roll, pitch, and remaining ground clearance of the vehicle, 3) these values are remapped between 0 (non-traversable) and 1 (traversable) and then thresholded, using a sigmoid function and the static performances of the vehicle (maximal frontal and side slope, ground clearance). The final traversability value assigned to the cell is the value of the least favorable of the three criteria. If one of the three criteria exceeds the robot’s limits, the cell is marked as non-traversable.

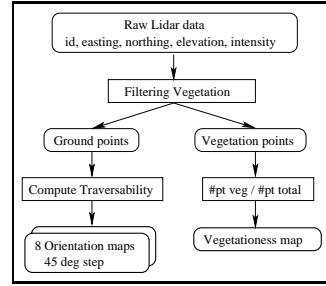


Figure 4: Flowchart of the computation of the traversability map

Planner To test our approach we implement a grid-based path planner to determine the path of least cost in our gridded map. The cost at each node is computed as follow :

$$C_{comb}(\theta) = \frac{1}{(1 - C_{trav.}(\theta))^2} + \frac{1}{(1 - C_{veg.})^2} \quad (1)$$

With $C_{trav.}(\theta)$ the directional traversability of the cell with θ the heading of the robot when it entered that cell ; $C_{veg.}$ is the *vegetationness* value of the cell; finally $C_{comb}(\theta)$ is the cell cost used to determine the path in this specific cell along that direction.

We tested the relative influence of the traversability and *vegetationness* map on the results produced by the path planner with the Yuma data set. We performed 3 different tests using 47 pairs of starting/ending points, selected randomly in the scene. We computed a path for each of them using three different sets of maps: 1) the 8 directional traversability maps and the *vegetationness* map, 2) one directional map and the *vegetationness* map, 3) the 8 directional traversability maps only. Each path produced, 141 total, has been evaluated visually using a high resolution aerial image (17cm) as ground truth. The table 3 summarized the results obtained. In all cases, a valid path is known to exist and a failure is recorded whenever a path cannot be generated from the cost map. The table contains in addition to the failure rate, the average length of the paths.

The importance of using the directionality map is clear from the experiments. The role of the vegetation map is not critical on the Yuma data set because of the nature of

Eval.	Trav.	Veg.	Failure rate	Av. length
#1	8 dir.	Yes	4.2%	528.9 m
#2	1 dir.	Yes	19.2 %	534.7 m
#3	8 dir.	No	4.2%	521.1 m

Table 3: Statistics on the relative influence of the cost map

the terrain: desert with tall bushes and trees with sparse vegetation. There is no continuous canopy tree as in the example presented in the next section.

5.2 APHill Data

Figure 7 presents an example of path obtained with our planner using the APHill data. The terrain is made of tall trees producing a continuous canopy cover of the ground. The density of foliage is such that sufficient ground points were sensed to reconstruct the terrain below the canopy. Figure 7-(a) is an aerial view of the scene with the initial 15 way points plotted in black. The total cumulative distance between each way points is 152 meters.

Figure 7-(b) represents the vegetation map. The path computed by the path planner is overlaid in black.

Figure 7-(c)/(d) are the traversability map computed respectively after filtering and without filtering the vegetation. The green areas are traversable and the dark red terrain is non-traversable. Points in blue are part of a water pond. The interest of the method is explicit on this example. During the demonstration a different path planner was used to produce a similar path passing through the same way points [2]. The robot actually navigated autonomously along this path, avoiding local obstacles (small obstacle, overhanging branches) not perceived in the aerial LADAR data.

6 Conclusion

In this article, we presented results on the use of aerial LADAR data for 3D terrain registration and path planning obtained during three different field demonstrations, using a ground mobile robot. We showed that because vegetation hides potential obstacles, masks the ground terrain features, and introduces artifacts which mislead point selection strategies, it is a major problem for mobile robot navigation in natural environment. We presented the methods used to filter vegetation both for ground and aerial LADAR data. Once the ground terrain is recovered, we have been able to produce ground-ground / air-ground 3D terrain registration for terrains in which registration would not be possible without vegetation filtering. We proposed to use ground terrain recovery below the canopy for path planning. We demonstrated with an autonomous ground robot that this new approach is viable.

Although encouraging, these results are still limited,


and work still needs to be done to further evaluate performance and to extend the approach to reach operational performance. Key directions of future work include: systematic evaluation of our approach with correlation to terrain ground truth, making the vegetation filtering method invariant to the terrain slope, and improvement of path planning by taking into account the clearance between the terrain and the bottom of the canopy.

Acknowledgments

This project was funded by DARPA under the PerceptOR program, under subcontract to General Dynamics Robotic Systems. This work would not have been possible without the help of William Klarquist from PerceptTEK.

References


- [1] E.P. Baltsavias, Airborne laser scanning: existing systems and firms and other resources *ISPR Journal of Photogrammetry & Remote Sensing*, 1999, vol 54
- [2] S. Balakirsky and A. Lacaze World modeling and behavior generation for autonomous ground vehicle, *Int. Conf. on Robotics and Automation*, 2000
- [3] D. Coombs, K. Murphy, A. Lacaze and S. Legowik, Driving Autonomously Offroad up to 36 km/hr *IEEE Intelligent Vehicles Symposium*, 2000
- [4] W. Eckstein and O. Munkelt Extrating objects from digital terrain models *Remote Sensing and Reconstruction for Three-Dimensional Objects and scenes*, SPIE Proceedings vol 2572, 1995
- [5] S.O. Elberink and H.G. Mass, The use of anisotropic height texture measures for the segmentation of airborne laser scanner data *Int. Archives of Photogram. and Remote Sensing*, vol XXXIII
- [6] J. Hyypya, O. Kelle, M. Lehtikoinen and M. Inkinen A Segmentation-based method to retrieve stem volume estimates from 3-D tree height models produce by laser scanners, *IEEE Trans. on Geoscience and Remote Sensing*, Vol 39, No 5, May 2000
- [7] D. F. Huber and M. Hebert, A New Approach to 3-D Terrain Mapping, *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, 1999
- [8] H. Jिंगgang, A.B. Lee and D. Mumford, Statistics of range images *IEEE Int. Conf. on Computer Vision and Pattern Recognition*, 2000
- [9] A. Johnson, Surface Landmark Selection and Matching in Natural Terrain, *IEEE Int. Conf. on Computer Vision and Pattern Recognition*, 2000




(a) Aerial view of the scene with the robot trajectory




(b) Terrain picture



(c) Aerial LADAR data. Elevation is encoded in color




(d) Side view of the ground LADAR data. In white the terrain surface, in color the vegetation filtered




(e) Registration results. In grey the aerial data. In red the ground data. In both case the vegetation has been filtered

Figure 5: Terrain registration with vegetation filtering. Yuma data set.


- [10] A. Johnson, Spin-Images: A Representation for 3-D Surface Matching, *Ph.D Thesis, Carnegie Mellon University*, 1997
- [11] K. Krauss and N. Pfeifer, Determination of terrain models in wooded areas with airborne laser scanner data *ISPRS Journal of Photogrammetry & Remote Sensing*
- [12] M.A. Lefsky et al., Lidar Remote Sensing of the Canopy Structure and Biophysical Properties of Douglas-Fir/Western Hemlock Forests *Remote Sensing Environment*, vol 70, 1999
- [13] P. Lohmann, A. Koch & M. Shaeffer, Approaches to the filtering of laser scanner data *Int. Archives of Photogrammetry and Remote Sensing* vol XXXIII, 2000
- [14] J. Macedo, R. Manduchi and L. Matthies, Laser-Based Discrimination of Grass from Obstacles for Autonomous Navigation *Int. Symp. on Experimental Robotics*, 2000
- [15] B. Petzold, P. Reiss and W. Stossel, Laser scanning surveying and mapping agencies are using a new technique for the derivation of the digital terrain models, *ISPRS Jour. of Photogram. and Remote Sensing*, 2000
- [16] C. Rasmussen, Combining Laser Range, Color and Texture Cues for Autonomous Road Following, *IEEE Int. Conf. on Robotics and Automation*, 2002.
- [17] G. Sithole, Filtering of laser altimetry data using a slope adaptive filter *ISPRS workshop on Land Surface Mapping and Characterization using laser altimetry*, 2001
- [18] N. Vandapel and M. Hebert, 3D Rover Localization in Airborne LADAR Data *Int. Symp. on experimental Robotics*, 2002



(a) Far view. Elevation encoded in color.



(b) Close view. Elevation is encoded in color




(c) Side view of the two terrain surfaces (no vegetation) registered


Figure 6: Ground-Ground registration in a terrain highly cluttered by trees. Pickel Meadow.




(a) Aerial image of test area with the way points as black dots



(b) Vegetation map (color scale from green to red: no vegetation to highly vegetated. In black the path computed



(c) Traversability map with the path overlay. Vegetation has been filtered. Color scale, from green to red: traversable to non-traversable



(d) Traversability map. Vegetation has not been filtered. No path has been found. Color scale, from green to red: traversable to non-traversable

Figure 7: Path planning and vegetation filtering. APHill