Hazard Detection and Avoidance

By Melissa Rowland, ECE '21

Introduction

As humanity expands its capabilities to explore space, there is a need to land assets on foreign bodies autonomously. Autonomous landing is especially crucial for unmanned missions to Mars or asteroids but will also be incorporated into the next lunar human lander system. Improving the precision and success rate of landing systems will extend capabilities to land in areas of scientific interest on foreign bodies that may have been too risky to visit in the past. An important component of any autonomous landing is the hazard detection and avoidance (HDA) system to ensure the safety of the vehicle in landing [1]. The landing system uses HDA in the later stages of the flight to best support the mission goal of landing safely in a region of scientific interest. Hazard detection algorithms generally take in maps of the surface features and determine hazardous landing spots based on surface roughness, surface angle, or object size [1]. Hazard avoidance algorithms take various approaches to determine the optimal landing site based on the hazards and surface features. Team Shamrock's goal is to simulate a navigation and HDA system for an asteroid landing in this senior capstone project. The HDA system will generate hazard maps from camera images and Lidar data, and choose safe landing sites that maximize the distance to the nearest hazard.

Background

Hazard Detection

The goal of hazard detection is to output a hazard map of the surface that shows where the hazards are located and where it is safe to land. The decision of which surface characteristics must be mapped from sensor data is driven by landing conditions that are of concern to the specific mission. For instance, many lunar missions are concerned with slope, surface roughness, and boulder size [1]. Figure 1 shows an example of hazards that may be identified during a lunar descent. A hazard map for each characteristic is generated by identifying which regions have values above the safety threshold [1]. An overall hazard map is generated from the map for each characteristic by taking the logical OR of them [1]. So, each spot in the final hazard map that is identified as safe is under the necessary threshold for all characteristics.

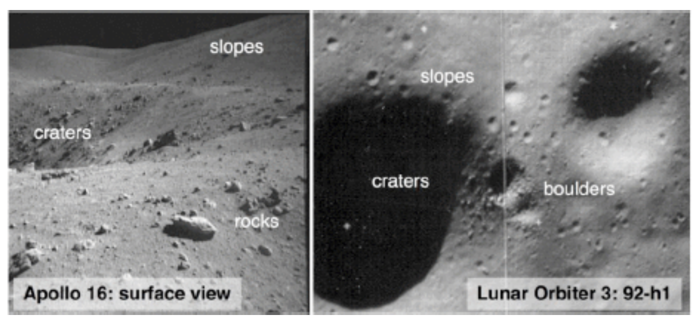

Figure 1. Hazards on the lunar surface [1]

In our simulation, Team Shamrock is concerned with the surface characteristic of rock size. A combination of passive and active imaging of the surface is used to make maps of hazardous rock locations. Passive imaging techniques use cameras to take images of the surface. Hazardous rocks are detected in the image using computer vision and image processing techniques. For more information, see the tech note on computer vision. LiDAR is used for active imaging of the surface to measure rock height directly. See the tech note on LiDAR for more information. Rocks that are determined to be above the threshold size will be marked as hazardous. The hazard maps output by the different sensors are combined using a logical OR. The combination of active and passive imaging makes the system more robust since the active system may be able to detect hazards obscured by shadows that will not be visible in the camera images. Additionally, for the passive techniques, the combination of computer vision and image processing modules developed by the team complement each other. The computer vision module is more successful at detecting smaller rocks while the image processing module is more efficient at detecting larger rocks.

Hazard Avoidance

The hazard map is the input to the hazard avoidance module. This module determines whether the planned landing site is safe, or if alternate landing sites must be chosen to avoid a hazard. The module outputs a list of possible landing sites, ordered from safest to less safe. This list is passed to a GN&C system, which will evaluate the sites along with other constraints such as remaining fuel levels and distance to target landing site to select the best landing site at the system level [1].

Hazard Avoidance Approaches

Maximize Distance to Nearest Hazard

A simple method to choose the safest landing site is to choose the site that maximizes the distance to the nearest hazard [1]. This site is chosen so that the lander will have the lowest probability of landing on a known hazard. First, the hazard map must be processed to create a Distance to Nearest Hazard (DTNH) map. An image processing technique called a grassfire transform is applied, which computes distances from pixels to the border of a region [1]. This technique allows the program to compute the distance from each pixel in the image of the landing area to the border of a hazardous region. The algorithm will then choose the landing site with the largest DTNH value as the optimal landing site [1]. After choosing the ideal site, to continue generating a ranked list of possible landing sites, the ideal site will have its DTNH value set to 0 and the site with the next highest value will be chosen [1]. This process is repeated until a list of desired length is generated. A tie breaker for choosing sites with the same DTNH value is to choose the site closest to the original desired site [1].

Figure 2 shows the output of Team Shamrock's hazard avoidance module using this approach. A hazard map with randomly placed rocks of various sizes was generated. The distance of each pixel to a hazard or the border of the image was computed. (The region just outside the border of the image is considered a hazard since it is unknown.) The ideal landing site shown as a red cross is the location with the maximum distance to any hazard.

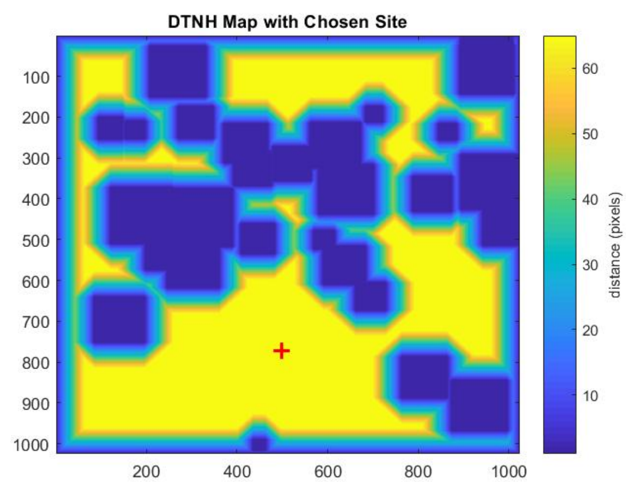

Figure 2. DTNH map with simulated random hazards, overlaid with ideal landing site (red cross).

Optimize Characteristics of Landing Site

The approach described above is simple to implement yet it assumes that the probability of a safe landing only depends on the probability of landing in a hazardous area. This is based on a binary assumption that a site is safe if the measures of its characteristics are below the safety thresholds and unsafe if the measures are above them. However, this approach ignores the fact that the measures of the characteristics can take on any value and may increase the probability of unsafe landing before the threshold is reached. A more nuanced approach is to generate a cost function for the map of the landing area that considers the values of each characteristic directly.

For an example of how to implement a cost function, let us consider a situation in which the characteristics of interest are surface roughness and landing angle [2]. First, the map is split into grid cells indexed with row r and column c. Any cell that is hazardous (exceeds the safety threshold) based on surface roughness, angle, or object size is assigned a cost value of $C = 1$ [2]. If any cell is within a distance of L (footprint of the lander) from a cell with cost $C = 1$, then it is also assigned $C = 1$ [2]. A cell with $C(r,c) = 1$ is considered unsafe. So far, this is the same way that the hazard map is generated for the distance to nearest hazard algorithm. The cost of the remaining grid cells are assigned the normalized product of landing incidence angle and roughness:

$$C(r,c) = [R(r,c) * A(r,c)] / (R_{max} * A_{max})$$

where $R(r,c)$ is the roughness at (r,c), $A(r,c)$ is the surface angle at (r,c), R_{max} is the maximum surface roughness, and A_{max} is the maximum landing angle [2].

The cost map is then smoothed since the selected site should be near regions of similarly low cost. This is achieved by setting the cost of a grid cell to be the average of all costs in a square neighborhood of size L centered at that cell [2]. The best safe landing site is selected to be the grid cell with the minimal value for the smoothed cost function. An ordered list of safe landing sites may be generated by assigning the cost of the previously chosen site and its neighbors to $C = 1$, and once again choosing the cell with the minimal value [2].

Figure 3 shows hazard maps of surface roughness and incidence angle generated from a smoothed terrain map. It also shows the cost function over the terrain computed from these maps.

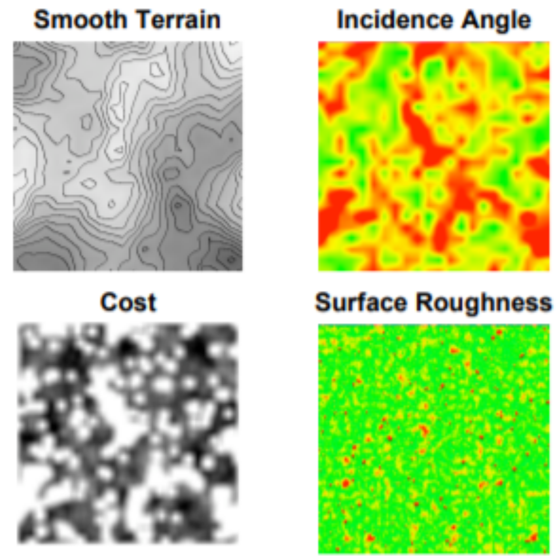

Figure 3. Cost Function for Mars Landing Surface [2]

System Level Considerations

The hazard detection and avoidance (HDA) modules play a key role in the success of the mission by ensuring safety, but there are other system level factors that must be considered as well. The choice of when and how to use HDA is determined by how it will support the whole system.

Target Landing Site Selection

The target landing site for the mission is planned before the landing attempt. Based on a priori knowledge, the target is chosen with a balance of safety and scientific goals in mind. For instance, on the NASA mission to the asteroid Bennu, the four key properties of the target sites were deliverability, safety, sampleability, and scientific value [4]. The first two properties directly related to the ability of the lander to contact the surface without damage that will cause failure of the system. The second two properties related to the overall mission goal of collecting carbonaceous regolith from the surface to study on earth. During descent, the HDA modules will run as the lander approaches the target site to detect hazards that were unable to be seen previously.

Efficient Use of HDA

There are costs to running the hazard detection and avoidance modules. The algorithms require time and compute power to run. LiDAR also consumes significant power as an active sensor [2]. To maximize useful information from HDA and to minimize wasted energy, an efficient approach is to utilize a relay autonomous hazard detection and avoidance scheme [3]. In the relay, a coarse hazard detection is performed at a far height from the surface (about 2km) and a fine hazard detection is performed much closer (about 100m above the surface) [3]. This method aims to avoid an obvious, large scale hazard in the coarse search [3]. This reduces the risk that obstacles detected in the fine hazard detection are too difficult to avoid and reduces the amount of fuel needed to maneuver around them [3]. The fine hazard detection complements the rougher search by detecting smaller hazards that are too detailed to see from far away but still compromise the safety of the system [3]. Fortunately, since they are smaller, it is easier to maneuver around them in the later stages of landing.

Conclusion

Team Shamrock was able to successfully implement a hazard detection and avoidance system for an asteroid landing simulation. Hazard maps were generated using LiDAR and camera images to determine the location of rocks that are large enough to be hazardous during landing. These maps were combined using a logical OR so that all hazards detected by the various sensors could be avoided. The hazard avoidance module used an algorithm that maximizes the distance to the nearest hazard to choose the best landing sites.

Future improvements to this project would include extracting more data about surface characteristics from the images to create maps of surface roughness and landing angle. This data could be inputted into a refined hazard avoidance module that takes a more holistic approach to landing site selection by assigning a cost function to the possible landing area, optimizing the characteristics of the landing site. An additional improvement would be to simulate other system level considerations in the landing site choice such as scientific value and resource constraints.

References

1. A. E. Johnson, A. Huertas, R. A. Werner and J. F. Montgomery, "Analysis of On-Board Hazard Detection and Avoidance for Safe Lunar Landing," *2008 IEEE Aerospace Conference*, Big Sky, MT, 2008, pp. 1-9, doi: 10.1109/AERO.2008.4526301.

2. A.E. Johnson, A. Klumpp, J. Collier, and A. Wolf. "Lidar-based Hazard Avoidance for Safe Landing on Mars," *NASA Jet Propulsion Laboratory,* California Institute of Technology, 1-38.

3. Q. Jiang, S. Li, & T. Tao. (2016). "Innovative hazard detection and avoidance strategy for autonomous safe planetary landing." *Acta Astronautica, 126*, 66-76, doi: 10.1016/j.actaastro.2016.02.028.

4. H. L. Enos , A. T. Polit, D.S. Lauretta , P. Antreasian, C.A. Bennett , K. Berry, E. B. Bierhaus, R. Burns, H. C. Connolly, Jr., B.E. Clark, M.K. Crombie, D.N. DellaGiustina, S. Freund, D. Lorenz, M. Houghton, M. Moreau, R. Mink, M.C. Nolan, & K.J. Walsh. "OSIRIS-REX's SEARCH FOR A SAMPLE SITE: SELECTING THE PRIME (NIGHTINGALE) AND BACKUP (OSPREY) SITES ON ASTEROID (101955) BENNU." *51st Lunar and Planetary Science Conference,* 2020, 1-2.