Finding Travel Routes for Off-Road Ground Vehicles Using Satellite Imagery

BLAKE GARWOOD

ABSTRACT

One pressing issue for unmanned and self-driving vehicles is navigating off-road terrain. It is hard to get a clear view of the area around the vehicle, making it even more difficult to decide which direction to travel in. This issue plagues researchers looking for locations to land rovers on other planets. We will create an AI that, using satellite imagery, can find a safe path through off-road terrain; both on earth, and elsewhere. The AI model will predict areas of the image that are "travelable," and a pathfinding algorithm will be used to link these areas together and find viable paths. The paths produced can be connected to form a network of travelable land. This network of travelable land can be used to find routes from one node to another. In the case of Mars, identifying nodes with many different paths will identify potentially good places to land rover missions, or even settler missions. The training set will be constructed from known areas that have been traveled by rovers such as Curiosity and Pathfinder. Mosaics of these areas have been produced with a spatial resolution of about 25 centimeters/pixel. Unfortunately, most of Mars, like many other planets, does not have satellite imagery available with spatial resolution better than about 20 meters/pixel. This solution has a high potential for technology transfer and has uses both on other planets, as well as with autonomous and self-driving vehicles here on Earth.

I. INTRODUCTION

One pressing issue for vehicles is navigating off-road terrain. On Earth, it is hard to get a clear view of the area around the vehicle, making it even more difficult to decide which direction to travel in. Elsewhere, this issue plagues researchers looking for locations to land rovers on other planets. We will create an AI that, using satellite imagery, can find a safe path through off road terrain.

Our plan is to use existing path data from sources such as the Curiosity Rover to train an AI to find clear paths using satellite imagery. This problem has a high potential for technology transfer, and can be used both on other planets, as well as with autonomous vehicles here on Earth.

This work is important because it can be expanded to a wide range of applications. It will add capabilities to unmanned or self-driving ground vehicles. It will also identify traveling paths for rovers on other planets. There could also be many other relevant applications to this AI.

II. BACKGROUND

Selecting a landing site on another planet is a hard and important problem. On top of trying to find a site with a lot of scientific opportunities, we must find sites with clear ground. Rovers do not fare well on surfaces with many small rocks, or surfaces with great slopes. Figure 1 shows the significant damage the Curiosity Rover's wheels have incurred from driving over rocks.

Figure 1: Damage shown on Curiosity rover wheels due to small rocks. [3]

We have very limited data available about the surface of proposed landing sites. One of our main assets in selecting such sites is the Mars Reconnaissance Orbiter. This satellite orbits Mars, and scans images of the surface with the Compact Reconnaissance Imaging Spectrometer for Mars (CRISM).

The images taken by the CRISM spectrometer are randomly sampled all over the surface of Mars [1] . The only problem with these images is that they only are around 15-20 meters/pixel. Consequently, these images are only able to provide a general picture of the land, but there is no way to see small rocks with these images.

A technique commonly used when finalizing the decision on a landing site is the production of what is called an orthophoto. This is formed by using image processing techniques to merge many overlapping stereo images. These images are timely and expensive to produce because of the amount of separate scans required to produce them. An example orthophoto is shown in Figure 2.

Figure 2: Curiosity Rover landing site orthophoto. Northwestern Gale Crater. Resolution = 25 centimeters/pixel. [2]

The expensive nature of producing images where we can see small obstructions gives us motivation to look elsewhere for solutions. This paper will discuss using AI to determine travelable vs. untravelable land by employing spectral unmixing of hyperspectral images.

III. RELATED WORK

There is a fair amount of related work with image pathfinding. Most work involves using aerial images, and image processing techniques to achieve a binary classification of travelable and untravelable land. From here, common pathfinding algorithms are used to traverse the resultant graph produced by the binary image.

Comparison of Aerial Imagery and Satellite Imagery for Autonomous Vehicle Path Planning

Robert Hudjakov and Mart Tamre compare their path planning algorithms between satellite imagery and aerial imagery. Their method is based on a convolutional neural network, which takes an image input and classifies it into a few groups. The output is an image which only has four different possible values: houses, grass, debris, and roads.

Splitting the image into these different groups allows the road portion of the image to be turned into a binary mask. This road matrix can then be used as the basis for a pathfinding algorithm. From this point, the solution becomes a relatively simple pathfinding problem. [5]

Cross-Country Path Finding using PSO and BDO

This paper is focused on finding paths where there are no roads. However, it still uses a similar method to the previous paper as far as finding paths which can be travelled. The paper uses Particle Swarm Optimization (PSO) to compute a threshold value, and subsequently threshold the input image into two classes: clear vs. obstructed land. After refining this image using various image processing techniques, a path can be planned using Biogeography Based Optimization (BBO).

We can see that this pathfinding technique works in a similar way to the previous. First it distinguishes the travelable land from the untravelable land, and then it uses some pathfinding algorithm to find paths through this travelable land. [4]

IV. MODEL OF SOLUTION

The model is based around two main functions: the Sobel Gradient, and the A* Pathfinding Algorithm. The high-level flow of the program can be seen in Figure 5. Many supportive functions are used to help tie together the gradient and pathfinding functions. There are also some functions used to help achieve better results from the gradient calculation step.

Before taking the gradient, local noise is reduced by smoothing the image. This smoothing process decreases the amount of noise in smooth, flat land areas of the images. Smoothing is executed using Non-Local Means Denoising [6]. The difference is slight to the human eye, but can be seen in the zoomed-in image comparison in Figure 3.

The gradient is a core function of this solution because it is the way the program distinguishes between travelable and untravelable land. The gradient is found using the sobel operator. This is a convolutional method which calculates the gradient based on the derivative of the image. The derivative is calculated in the x and y directions independently, and then combined together by simply adding the resultant matrices. Finally, a binary image is created by thresholding the values at each individual pixel. Anything below the threshold value will become “travelable” (value of 0) in the binary image.

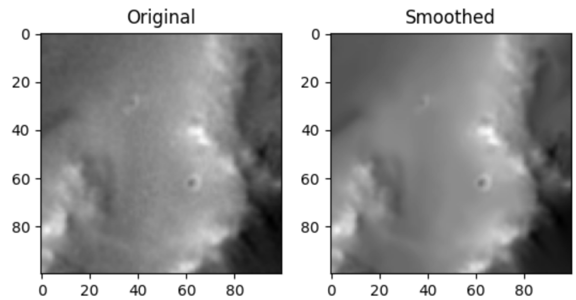

Figure 3: Image before and after “smoothing” process. Zoomed in for closer inspection.

After calculating the gradient, simple image processing techniques are used to enhance the produced binary image. Dilation and erosion help join together fragmented areas of untravelable land. Dilation grows areas that are a ‘1’ in the image by ‘OR’ing a 3 by 3 matrix of ones across all ‘1’ values in the image. Erosion does the exact opposite, shrinking the area. The net effect is that most isolated areas in close proximity to each other are joined together, theoretically resulting in no external boundary change. This can be seen in Figure 4.

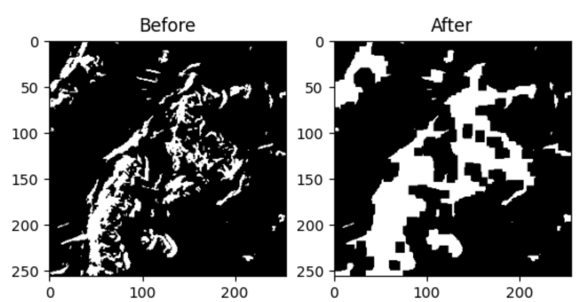

Figure 4: Before and after enhancing binary image.

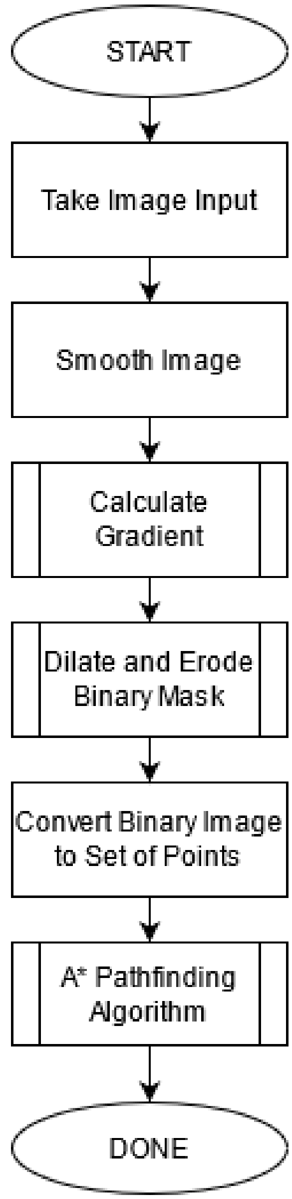

Figure 5: High-level routine for satellite pathfinding code.

The A* Pathfinding Algorithm is being used to traverse the resultant graph and find the shortest path between given points. This algorithm iteratively extends the current tree of routes by calculating the projected lowest cost of extending the path to the endpoint. This ensures a relatively quick runtime. To implement the A* Pathfinding Algorithm in our project code, we used open source python files from programmer Atsushi Sakai [7].

V. RESULTS

This algorithm was tested on 20 different 256x256 images of the Martian surface. The results were variable depending on the contrast of the image. Similar images to the calibrated test images fared well in testing. These images had their terrain features detected properly by the gradient algorithm. However, images that were higher or lower contrast had different results.

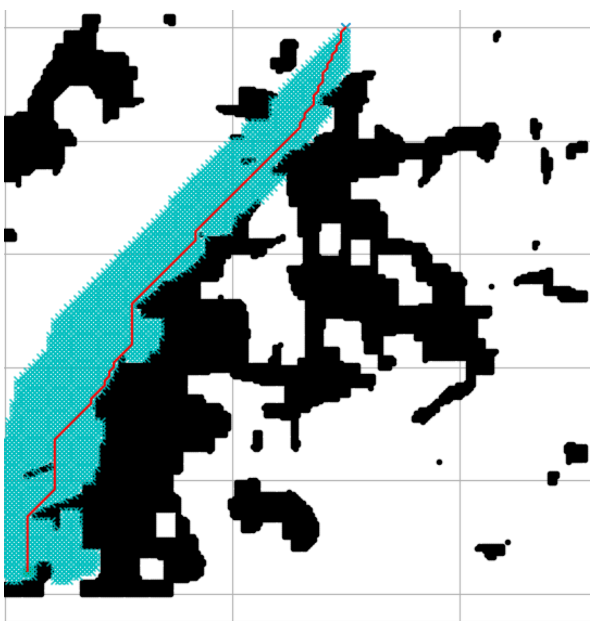

Figure 6: Visualization of A* Pathfinding algorithm. Blue pixels are iterated sites. Red line is the calculated best path.

The main variability in results relied on honing in the gradient threshold to form the binary mask. This was originally done through iterative testing. Values were tweaked and then the gradient function was re-tested multiple times. This helped achieve accurate results. Before expanding testing to multiple images, this process worked well, unfortunately, this process must be re-done when using new images with contrasts than the original image.

One issue with the resulting binary image is the open spaces the erosion has created. A few distinct rectangular holes can be seen in the middle of the mountain range in Figure 6. One improvement that could be made is filling in these holes, so no start or end node could be accidentally placed there.

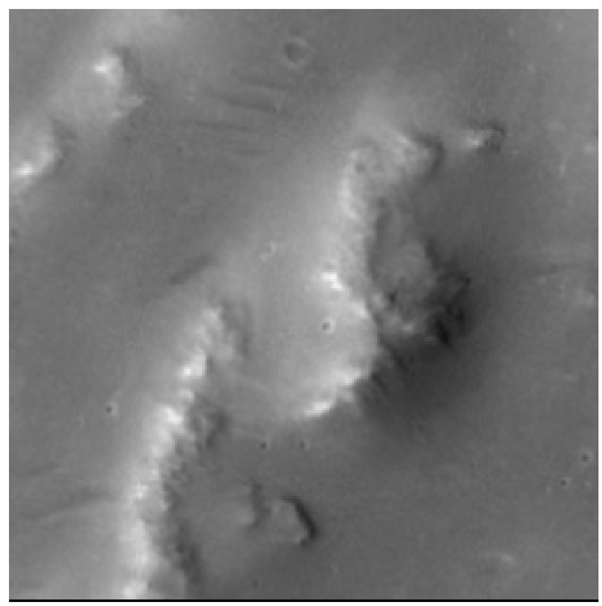

Figure 7: Example of original image before preprocessing. Same image used for visualization in Figure 6.

VI. FURTHER WORK

Further work will be carried out on this program to fix some of the issues, as well as add greater capability. The first work that will be done is to try to normalize the contrast of the input image to achieve more reliable results with the gradient calculation function.

We can attempt to normalize the contrast by generating a color balance map of an original test image. This will be an image for which the values of the gradient function as well as the dilation and erosion functions are calibrated. In the input, a function will be created which changes the contrast by re-scaling the brightest white and darkest black pixels. The rest of the pixels will be scaled along this new range of values.

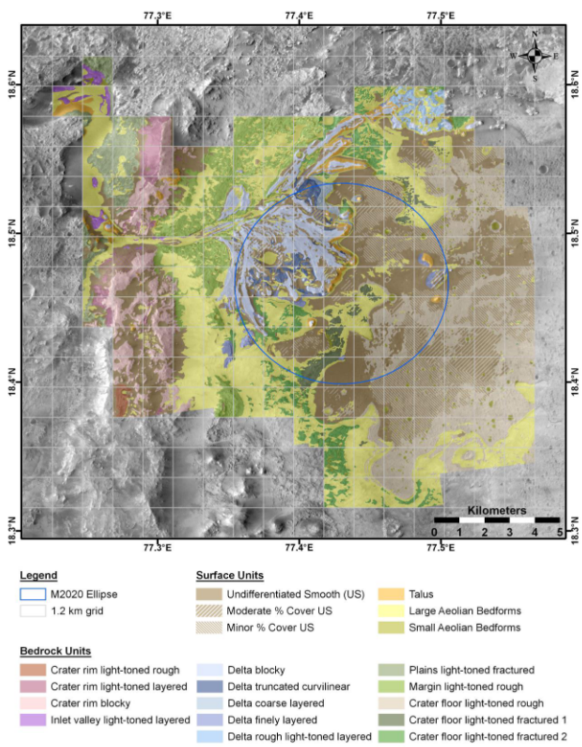

Figure 8: Photo-geologic map of Jezero Crater. The blue circle is the main mission area for the Perseverance Rover.

We will also increase the capability of this program by using hyperspectral signatures of low-spatial resolution images for pixel-by-pixel detection of rocky "untravelable" terrain. These rocky terrain pixels will be added to the binary image of untravelable land.

This will be carried out by overlaying low-resolution images onto produced orthophotos with corresponding maps. Photo-geologic maps are produced before each NASA rover landing. The map featured in Figure 8 is an example of one. Because they contain information about the surface material composition, these maps can be quickly transitioned into travelable/untravelable binary classification maps.

Geographical features can be used to align these maps with the lower-resolution satellite images. This will give us an accurate training set to use to train an artificial neural network which hyperspectral signals can be traveled across.

Finally, this produced neural network can be used to classify the pixels in new images. The result will be a second masked image which can be combined with the first using a simple 'OR' operation. This will allow the program to avoid the rocky terrain type previously mentioned.

VII. REFERENCES

[1] *CRISM Data Products*, crism-map.jhuapl.edu/.

[2] “Mars MSL Gale Merged Orthophoto Mosaic 25cm v3: USGS Astrogeology Science Center.” *Mars MSL Gale Merged Orthophoto Mosaic 25cm v3 | USGS Astrogeology Science Center*, astrogeology.usgs.gov/search/map/Mars/MarsScienceLaboratory/Mosaics/MSL_Gale_Orthophoto_Mosaic_10m_v3.

[3] “Wheel Treads Break on Curiosity Rover.” *SpaceFlight Insider*, 8 June 2017, www.spaceflightinsider.com/missions/solar-system/wheel-treads-break-curiosity-rover/.

[4] Kundra, Harish, and Monica Sood. "Cross-country path finding using hybrid approach of PSO and BBO." *International Journal of Computer Applications* 7.6 (2010): 15-19.

[5] Hudjakov, Robert, and Mart Tamre. "Comparison of Aerial Imagery and Satellite Imagery for Autonomous Vehicle Path Planning." *8th International DAAAM Baltic Conference*. 2012.

[6] Antoni Buades, Bartomeu Coll, and Jean-Michael Morel, *Non-Local Means Denoising,* Image Processing On Line, 1 (2011), pp. 208-212. https://doi.org/10.5201/ipol.2011.bcm_nlm

[7] atsushisakai.github.io/pythonrobotics/