



Mars 2020 Surface Mission Performance Analysis: Part 2. Surface Traversability

Masahiro Ono¹, Matthew Heverly², Brandon Rothrock³, Takuto Ishimatsu⁴, Eduardo Almeida⁵, Fred Calef⁶, Tariq Soliman⁷, Nathan Williams⁸, Hallie Gengl⁹, Austin Nicholas¹⁰, Erisa K. Stilley¹¹, Kyohei Otsu¹², Marshall Trautman¹³, Robert Lange¹⁴, and Sarah Milkovich¹⁵

Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, 91106

Abstract – The Mars 2020 Rover Mission (M2020) is characterized by demanding requirement on the distance and time for traveling between scientific Regions Of Interest (ROIs). As a result, surface traversability is one of the major driving factors for the landing site selection of M2020. With the newly developed Mars Terrain Traversability analysis Tools (MTTT), we performed traversability analysis of the eight candidate landing sites with an unprecedented granularity. This paper describes the MTTT analysis capabilities, as well as how the MTTT capabilities were used to downselect from eight to three candidate landing sites for further evaluation.

Table of Contents

Mars 2020 Surface Mission Performance Analysis: Part 2. Surface Traversability	1
I. Introduction	2
A. Landing Site Candidates	2
B. Baseline Reference Scenario	3
II. Analysis Method	4
A. Overview	4
B. Map	4
C. Terrain Classification	5
D. Rock Abundance	6
E. Slope	7
F. Mobility Model	7
G. Time-Optimal Route Planning	8
H. Monte Carlo Analysis	9

¹ Research Technologist, Mobility and Robotics Systems Section, ono@jpl.nasa.gov, non-member
² Supervisor for Mechanisms and Mobility, Mechanical Engineering Section, mheverly@jpl.nasa.gov, non-member
³ Research Technologist, Mobility and Robotics Systems Section, Brandon.Rothrock@jpl.nasa.gov, non-member
⁴ Systems Engineer, Project Systems Engineering and Formulation Section, takuto.ishimatsu@jpl.nasa.gov, non-member
⁵ Former member of JPL/Caltech
⁶ Data Scientist, Science Data Understanding Group, Fred.Calef@jpl.nasa.gov, non-member
⁷ Former intern student of JPL/Caltech
⁸ Postdoctoral Scholar, Geophysics and Planetary Geosciences Section, nathan.r.williams@jpl.nasa.gov, non-member
⁹ Data Visualization Developer, Instrument Software and Science Data Systems, hallie.e.gengl@jpl.nasa.gov, non-member
¹⁰ Systems Engineer, Project Systems Engineering and Formulation Section, Austin.K.Nicholas@jpl.nasa.gov, non-member
¹¹ EDL Systems Engineer, Erisa.K.Hines@jpl.nasa.gov, non-member
¹² Robotics Technologist, Mobility and Robotic Systems Section, otsu@jpl.nasa.gov, non-member
¹³ Systems Engineer, Software Systems Engineering Section, marshall.r.trautman@jpl.nasa.gov, non-member
¹⁴ Systems Engineer, Mission Systems Engineering Section, Robert.D.Lange@jpl.nasa.gov, non-member
¹⁵ Science Systems Engineer, 394B, sarah.m.milkovich@jpl.nasa.gov, non-member

III. Analysis Results	9
A. Results Overview	9
B. Interpretation	11
C. Qualitative Assessment	13
IV. Conclusion	14
Acknowledgments	14
References	14

I. Introduction

THIS paper presents the analysis approach and results on the surface traversability of the eight candidate landing sites considered in the Third Landing Site Workshop for the Mars 2020 Rover (M2020) mission. Achieving the goals of the M2020 mission is even more contingent on mobility than Mars Science Laboratory (MSL). MSL's science goal is incremental, meaning that the more the rover drives the greater science return you get. In contrast, the goal of M2020 is somewhat binary, meaning that mission success cannot be declared unless the rover successfully visits two geographically dispersed Regions Of Interest (ROIs) and complete sample collection. As a result, in the ongoing landing site selection process for M2020, a greater emphasis is placed on traversability analysis than any of the past Mars missions. We developed a suite of traversability analysis capabilities called *Mars Terrain Traversability-analysis Tools (MTTT)*. In this paper we will show how the MTTT analysis capabilities were used to downselect from eight to three candidate landing sites for further evaluation.

A. Landing Site Candidates

Table 1 lists the eight landing sites considered in the Third Landing Site Workshop. In the rest of this paper, we call the landing sites by three-letter abbreviations shown in the table.

Table 1 Candidate landing sites considered in the M2020's Third Landing Site Workshop

Site name	Abbreviation	Latitude	Longitude
Columbia Hills (Gusev Crater)	CLH	14.5478°S	175.6255°E
Eberswalde	EBW	23.7749°S	33.5147°W
Holden Crater	HOL	26.6200°S	34.8713°W
Jezero Crater	JEZ	18.4386°N	77.5031°E
Mawrth	MAW	23.9685°N	19.0609°W
Northeast Syrtis Major	NES	17.8899°N	77.1599°E
Nili Fossae Trough	NIL	21.0297°N	74.3494°E
Southwest Melas Basin	SWM	9.8132°S	76.4679°W

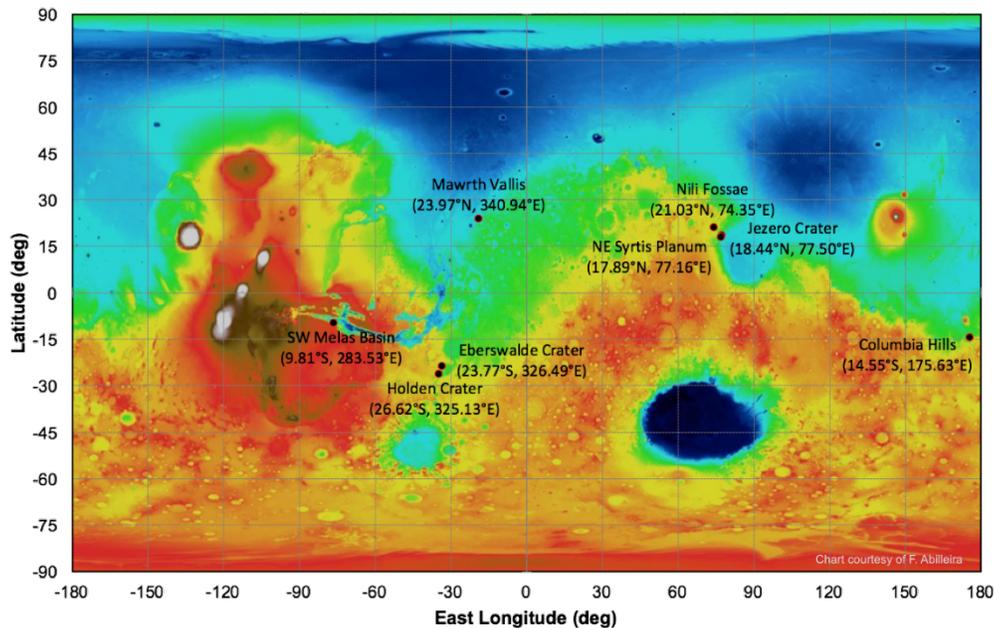


Figure 1 Distribution of the eight candidate landing sites for the Mars 2020 Mission

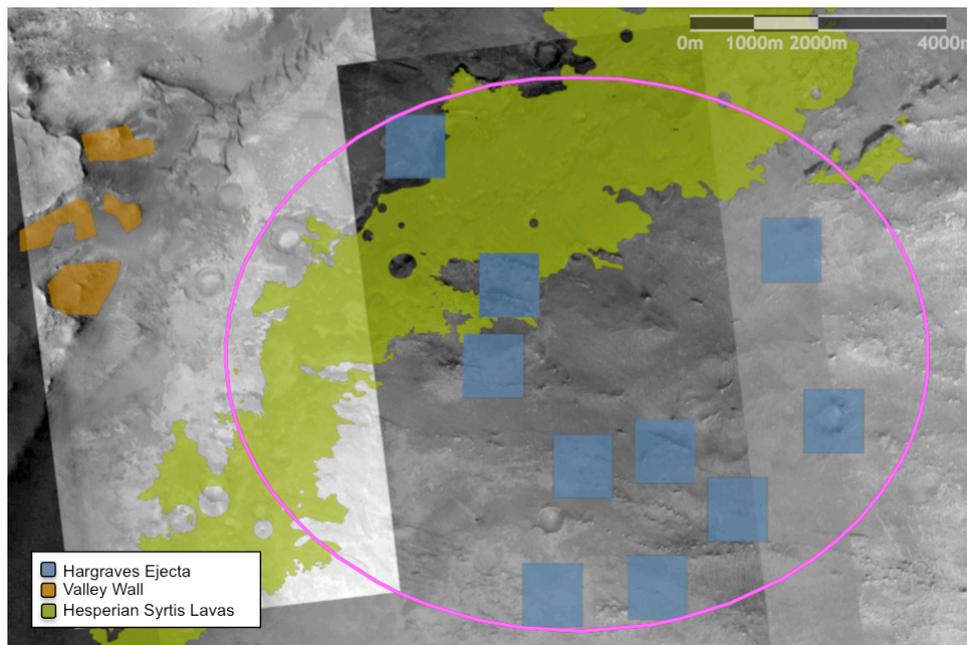
B. Baseline Reference Scenario

The notional mission scenario for the Mars 2020 rover is to traverse to two scientifically interesting and geologically diverse locations known as Regions of Interest (ROIs). The ROIs are approximately 1 km x 1km areas where the rover can carefully explore the most interesting geology in the region. Once the rover explores the first ROI, it will quickly drive to the second ROI, doing minimal science investigations along the way. The project constructed a Baseline Reference Scenario (BRS) as a generic representation of the mission. In this BRS, 85 sols and 12 km of driving distance is allocated to traversing between the landing location and the first ROI, and then traversing from the first ROI to the second ROI. This efficient traverse between ROIs allows the majority of the mission to be spent inside the ROI performing scientific exploration and collection of rock core samples for potential return to Earth.

To select a landing site, each of the candidate sites is evaluated against the BRS. Driving time and distance between the ROIs depends on the location of the ROIs in the landing ellipse, the presence of mobility hazards, and the traversability of the terrain. The analysis presented in this paper enables the project to evaluate the traversability between ROIs for each candidate landing site.

For each site, our team worked with the project scientists as well as the site proposers to specify candidate ROIs at each site. In addition to ROIs, candidate waypoints were also identified. These are areas of moderate scientific interest where the rover would perform an abbreviated exploration and sample caching campaign lasting only a few sols. For example, Figure 2 shows the ROIs and waypoint in Nili Fossae. Shown in blue (Hargraves Ejecta) and orange (Valley Wall) are the two ROIs. The waypoint can be anywhere in the area shown in green (Hesperian Syrtis Lavas). A typical mission would be to land somewhere in the ellipse, drive to one of the blue ROIs, then head west towards the orange ROI. When the path comes across the green unit, it stops and take samples (waypoint). The selection of path within ROI is driven by science and in-situ discovery and not suitable for a priori optimization. Thus, the analysis only considers optimization of the inter-ROI driving and the mission keeps an allocation of time for intra-ROI driving. Our path planner chooses a route that chooses the entry and departure points to/from ROIs such that the total driving time is minimized. In reality, the entry and departure points would be decided in consideration of the science campaign conducted in the ROIs. While ROI and waypoint are different concept in terms of operation, our route planner treats them equally because our analysis was only concerned with drives between ROIs/waypoints. Part 1 of this paper series explains in detail, the

scientific motivations for the ROIs and waypoints on each site. Details of BRS will be discussed in Lange et al Part 3.



II. Analysis Method

A. Overview

Our analysis focused on long-range traverses between ROI. For each landing site, we first created a several types of maps: slope map, rock abundance (CFA) map, and terrain map (Section II.B-E). These maps were used for assessing the ability to drive and estimating driving speed of the rover at a given location (Section II.F). This information is used by an optimal route planner to plan the minimum-time route from a landing point to the ROIs (Section II.G). For each site, we plan such an optimal route for 8,000 landing points, randomly sampled from the probability distribution provided by M2020 EDL (entry, descent, and guidance) analysis team (Section II.H). This Monte Carlo simulation produced the final output from the MTTT analysis, the probability distributions on driving time and distance for each site.

B. Map

We produced map documents to compile the suite of data products available for each landing site in a geographic information system (GIS), similar to previous landing site selection works for Pathfinder¹, Spirit and Opportunity², Curiosity³, and InSight⁴. Data types of interest for landing site safety analysis and selection include ellipses, ROIs, visible images, terrain classifications (Section II.C), rock abundance estimates (Section II.D), digital elevation models and slope measurements (Section II.E).

The Mars Orbiter Laser Altimeter (MOLA) planetocentric topography⁵ serves as the fundamental reference basemap for all surface data products. Visible image mosaics are sequentially co-registered to MOLA at increasing resolution. High Resolution Stereo Camera (HRSC) orthoimages at ~ 12.5 m/pixel⁶ are well co-registered to MOLA and serve as the base for other visible images. Distinctive small craters visible in both Mars Reconnaissance Orbiter Context Camera (CTX)⁷ and HRSC images are connected via tie-points at crater centers to co-register and thus geo-reference the CTX to a known base without bias from shadows. Craters are preferred for tie-pointing because their planform shape is always a circle in the same place regardless of shading, whereas a rock for example could cast a shadow in any direction. After CTX is geo-referenced, even smaller craters visible in both CTX and High Resolution Imaging

Science Experiment (HiRISE)⁸ images are tie-pointed to geo-reference the HiRISE images. For ortho-images where topographic distortions have been corrected, a first-order polynomial transformation is used, and images are typically placed as the highest layers in the image basemap mosaic. Other map-projected or non-map-projected images use a spline transformation to provide a rougher correction for such topographic distortions and are placed as lower layers in the image basemap mosaic. Higher resolution images are placed in higher layers for the basemap, such that 25 cm/pixel HiRISE is always overlying 6 m/pixel CTX which overlies 12.5 m/pixel HRSC. Image noise, jitter, saturation, poor illumination, and/or other artifact issues are also considered in layering images for the basemap to ensure that the highest quality images are seen. The tie-points used to geo-reference images to each other are also saved and can be applied with the same transformation to data products derived from those images (*e.g.*, terrain classifications, rock abundances, topography, slopes, and ROIs) and thus geo-reference them as well.

All 8 ellipses for the Third Landing Site Workshop have complete 6 m/pixel image coverage, although only Nili Fossae and Northeast Syrtis had topography-corrected orthoimages prior to the workshop. Only the Northeast Syrtis ellipse had full coverage with 25 cm/pixel HiRISE orthophotos as well; while Eberswalde, Holden, Jezero, Mawrth, Nili Fossae, and Southwest Melas had >75% orthoimage coverage; and West Columbia Hills had 40% orthoimage coverage. Although not as well corrected for topographic distortions, HiRISE map-projected images were also available to fill in additional gaps in HiRISE-resolution coverage: West Columbia Hills, Eberswalde, Jezero, and Nili Fossae were complete; Mawrth was 95% complete; Southwest Melas was 93% complete; and Holden was 79% complete. Since the workshop, Jezero now also has complete CTX and HiRISE orthoimage coverage, and Columbia Hills now has complete CTX orthoimage coverage with additional efforts currently underway towards improving HiRISE orthoimage coverage.

Prior to the Third Landing Site Workshop, we used a standardized projection for all landing sites of Equiarectangular Cylindrical with a center longitude of 0°E and center latitude of 0°N (also called Plate Carrée). This simple projection for low latitudes provides ~5% accuracy in distances from distortions, which was sufficient for the recent workshop and this work. For the future Fourth Landing Site Workshop, each site will use a customized projection to further reduce geographic distortions to <1%. Image-style products were exported as geotiffs, while vector products were saved as shapefiles.

C. Terrain Classification

A lesson learned from the operation of the Curiosity rover is that traversability and driving speed is significantly influenced by terrain type (*e.g.*, sand, regolith, bedrock, etc.) We identified 17 terrain classes relevant to traversability that exhibit unique surface morphology or visually distinctive patterns. These terrain classes are described in detail in⁹. All of our terrain analysis is performed on HiRISE imagery at the nominal native resolution of 25cm/px. Our automated terrain classifier is referred to as Soil Property and Object Classification (SPOC)⁹, and is based on a fully-convolutional neural network (FCN)¹⁰ trained from sparse annotations on HiRISE mosaics for each candidate landing site. The FCN architecture is attractive because it allows the network to be trained end-to-end from image pixels to label maps directly, without any post-processing or ad-hoc ensembles of models. Furthermore, the deep architecture of FCNs are capable of learning complex texture patterns and subtle geometric compatibilities between neighboring regions. Our implementation is a variant of “DeepLab”¹¹, which utilizes multiple dilated convolution layers and a single upsample layer. We refer the reader to¹¹ for details on the design and training of the network.

SPOC learns from a relatively small collection of human annotated examples, and applies the learned model to classify similar patterns across the remaining candidate landing site. Deep learning approaches such as this typically require a large number of training examples, which is both time consuming and expensive to collect. In contrast, our approach attempts to minimize the amount of training data required to produce a satisfactory terrain map for each candidate site individually. This is achieved by iteratively retraining the classifier and incrementally amending annotations to include corrections for misclassified regions. Performance of the classifier stabilizes after two or three iterations, typically resulting in less than 1% of the site being annotated, and approximately 90% average precision. Figure 3 below shows the examples given by human experts (left) and the classification by SPOC (right).

Annotations are created in a web-based multiscale map visualization tool that allows the user to overlay the current terrain map produced by the classifier with drawn vector-based regions for each terrain type. Training pairs are created by rasterizing the vector annotations with every HiRISE strip containing the region to create a label map. In many cases there are multiple HiRISE strips that overlap in the same region, providing multiple examples for the same label. Each HiRISE image is then locally contrast normalized by subtracting the mean and dividing by the local average

intensity. Training is computed on modern GPUs, but due to memory limitations the HiRISE image and label pairs are subdivided into tiles, discarding any tiles devoid of labels. Any remaining pixels that are unlabeled or flagged as NODATA are masked in the loss function, and do not contribute to the backpropagation gradients. A standard cross-entropy loss is used per-pixel. Once training is complete, classification is computed on individual tiles for each HiRISE image, and stitched into a single mosaic for each candidate landing site.

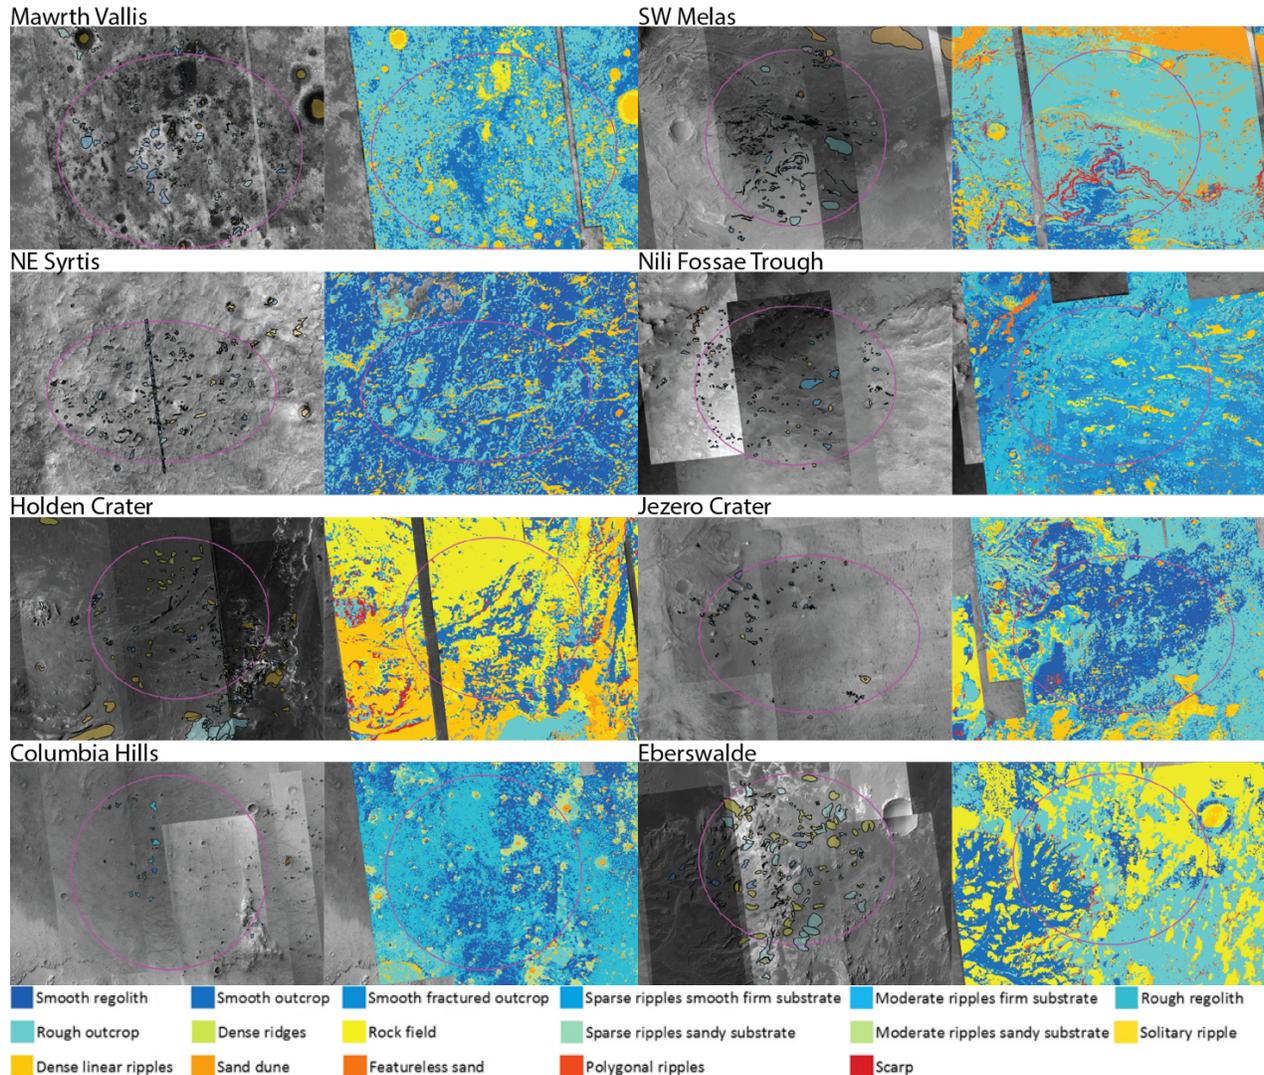


Figure 3 Terrain classification provided by human experts (left) and by SPOC (right) on all eight 2020 candidate landing sites. SPOC learns from the sparse human examples, and applies the same pattern to the entire site.

D. Rock Abundance

Besides terrain type, rock abundance is another major factor for traversability and driving speed. Rock abundance was determined by automatically detecting rocks on HiRISE non-map-projected imagery (~0.3m/pixel resolution), and fitting a model Cumulative Fractional Area (CFA) to filtered detected rocks. JPL's rock detection algorithm resolves the location and diameter of individual rocks by sharpening the image, enhancing the shadows, segmenting the shadows, fitting an ellipse to each shadow, and fitting a cylinder to each rock¹². Detections with 1.5-2.25m diameters were subset for CFA model fitting because the algorithm does not reliably detect rocks with diameters under 5 pixels due to resolution roll-off, and makes false positive detections on non-rock shadow-casting features, such as mounds, with diameters greater than 2.25m¹³. The CFA model

approximates the abundance of all rocks based on the size frequency distributions of diameters over the observation area¹⁴. A 30x30m grid was placed over each detection image, and CFA was fit to filtered rock detections within a 150x150m moving window.

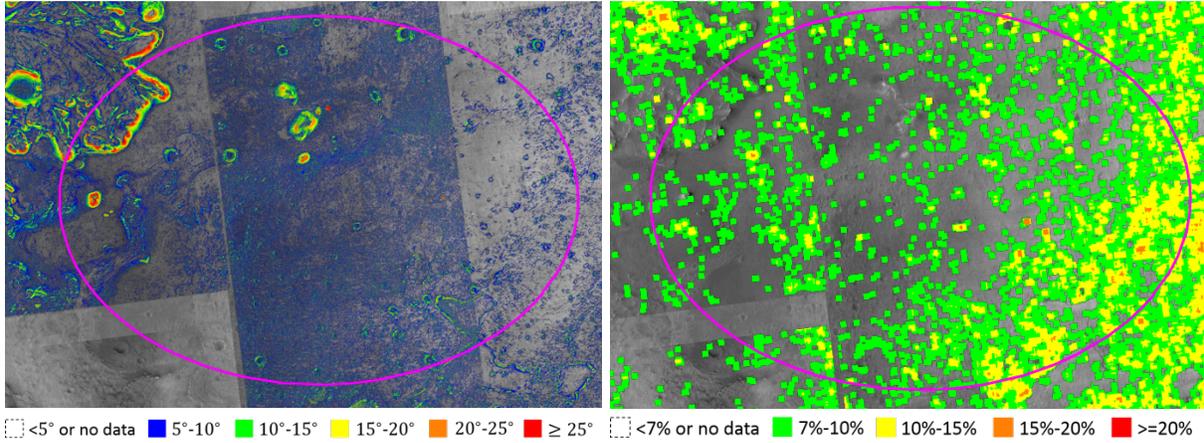


Figure 4 Slope (left) and CFA (right) maps of Jezero Crater.

E. Slope

Topography for landing sites is available as digital elevation models (DEMs) that were created using SOCET-SET software. Pairs of orbital 25 cm/pixel HiRISE images were taken with similar lighting conditions but different spacecraft orientations (emission angles) that produces a parallax effect where nearfield and farfield ground objects at different distances to the spacecraft appear shifted relative to one another, similar to human depth perception. Distinctive landmarks/features in both images are identified as tie-points, and the relative offsets between tie-points is measured and combined with camera viewing geometries to compute a mesh grid for terrain in overlapping images¹⁵. The resulting 1 m/pixel digital terrain is then vertically anchored to MOLA tracks to remove artificial tilts and approximate true elevations.

Slopes were calculated for available 1 m/pixel HiRISE DEMs using the Horn method¹⁶. At each elevation posting (step size of 1 m), a plane is fit by weighting its 8 surrounding/nearest neighbors (a 3 m by 3 m moving window excluding the center point):

$$Slope(x_i, y_j) = \tan^{-1} \left(\sqrt{\left[\frac{dz}{dx} \right]^2 + \left[\frac{dz}{dy} \right]^2} \right)$$

$$\left[\frac{dz}{dx} \right] = \frac{(z_{(i+1),(j+1)} + 2 * z_{(i+1),(j)} + z_{(i+1),(j-1)}) - (z_{(i-1),(j+1)} + 2 * z_{(i-1),(j)} + z_{(i-1),(j-1)})}{8 * x_{cellsize}}$$

$$\left[\frac{dz}{dy} \right] = \frac{(z_{(i-1),(j-1)} + 2 * z_{(i),(j-1)} + z_{(i+1),(j-1)}) - (z_{(i-1),(j+1)} + 2 * z_{(i),(j+1)} + z_{(i+1),(j+1)})}{8 * y_{cellsize}}$$

HiRISE DEMs and slope maps are available at all 8 ellipses but with varying degrees of coverage completeness. As of the Third Landing Site Workshop, the ellipse at Northeast Syrtis had full coverage; ellipses at Eberswalde, Holden, Jezero, Mawrth, Nili Fossae, and Southwest Melas had >75% coverage; and the ellipse for the west side of the Columbia Hills had 40% coverage. Since the workshop, work on the downselected sites has continued such that topography and slope coverage is now complete at Jezero, and work is ongoing to complete coverage at Columbia Hills.

F. Mobility Model

Based on the experience from *Curiosity* as well as our best knowledge of the Mars 2020 rover's mobility capability (e.g., autonomous navigation), we developed a mobility model that maps terrain type, CFA, and slope to expected traverse rate, as shown in Figure 5 below. The 17 terrain types are categorized into five classes. In each class, the expected rate is given as a 2-D function of CFA and slope. Note that shown in Figure 5 is the average driving rate over 2.5 hour driving time per each Sol (Martian day). The estimated drive rate also accounts for vehicle slip and path inefficiency (i.e., extra distance to avoid obstacles) which are also a function of terrain type, slope, and CFA. In some terrains we expect the rover to be capable of driving beyond where the human operators have imagery using the rover's onboard autonomous navigation. In other terrains, however, where the risks are too great (e.g. steep slopes or extremely rocky terrain) we expect the rover to drive only where humans can confirm it is safe, so the effective traverse rate is quite small (30 m/sol in the allocated 2.5 hours of drive time with 10% path inefficiency yields an effective rate of only 10.9 m/hr).

In order to capture the uncertainty of the traverse rate in certain terrains, we use two sets of driving rate models, one with conservative assumptions and the other with optimistic assumptions. These models are primarily based on uncertainty in how the autonomous navigation algorithm will work in rough and sandy terrain. For rough terrain, it is difficult to tell from orbit if paths exist that the rover will be able to find autonomously. If the rover is unable to find a safe path it will stop driving and rely on human operators to specify its path. In sandy terrain, the autonomous algorithm may not be capable of sufficiently sensing the terrain. Since the autonomous navigation depends on stereo imagery to build a terrain hazard map, if the sand lacks sufficient features, the algorithm won't be able to build a dense stereo map. In this case of featureless sand, human operators will also be necessary to specify the path. In both of these instances, optimistic assumptions make full use of the autonomous navigation capability and conservative assumptions assume that all motions are directed by human operators and limited to 30 m/sol on average.

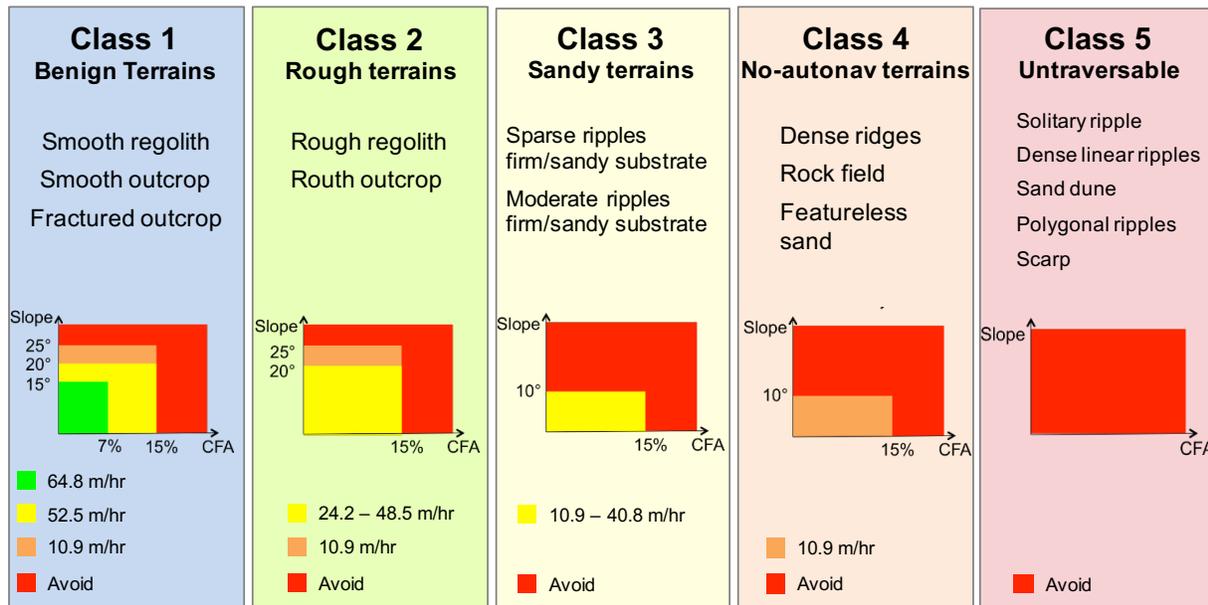


Figure 5 The driving rate model we used for the analysis

G. Time-Optimal Route Planning

The optimal route planner, which is built upon the Sequential Dijkstra Algorithm¹⁷, was used to find a path that starts from a given landing point and visits at least one point (i.e., sample point) in each of ROIs and minimizes the driving time. We first built a cost map, whose entry at each cell is the inverse of the estimated driving speed (i.e., time to drive a unit distance) derived in the previous section. The Sequential Dijkstra Algorithm computes a path that minimizes the total cost on the cost map, which is the minimum-time path.

The estimated driving speed has ranges for Classes 2 and 3. Accordingly, two cost maps were created: one for the conservative estimate and the other for the optimistic estimate. The resulting optimal routes are different. In general, as the conservative estimates result in higher cost on rough and sandy terrains, the optimal routes with the conservative estimate tend to stick to benign terrains by avoiding rough and sandy terrains, which in turn result in greater path length and driving time.

Figure 6 shows the visualization of paths with 100 landing points on Jezero (left) and Holden (right) with the conservative driving speed estimates. The color of the path represents the driving speed (green being fastest), while the thickness represents the frequency that the route is traveled. Interestingly, there are natural “highways,” which allows faster drive than the surrounding terrain. As a result, many paths converge to the “highways”. In the case of Holden (right), for example, the majority of the landing site is covered by moderate ripples, a terrain that the rover can only drive at reduced speed. The highways are on smooth outcrop terrains. A typical path starts on moderate ripples, drives to the nearest highway (orange portion), and then drives to ROIs by following the highways.

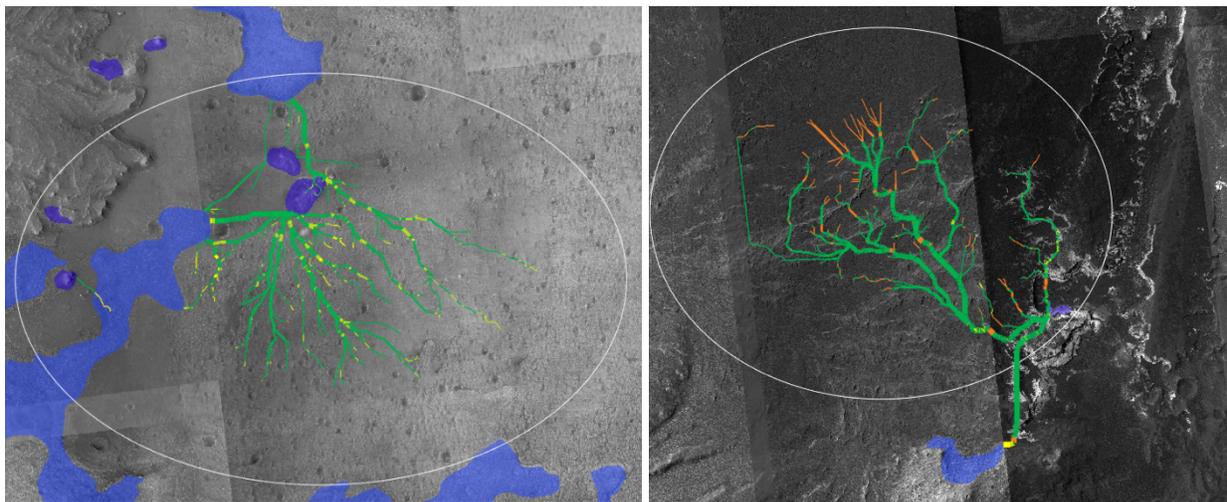


Figure 6 “Highway charts” for Jezero (left) and Holden (right), which show 100 paths. The color represents the driving speed (green being the fastest), while the thickness represents the frequency that the route is traveled. Two ROIs are shown in blue.

H. Monte Carlo Analysis

Monte Carlo simulations were performed to obtain the probability distribution of the driving distance and time. For each landing site, we ran the optimal route planner starting from 8,000 landing points with both conservative and optimistic driving speed estimates. The landing points were sampled from the probability distribution based on entry, descent, and landing (EDL) simulations with the meso-scale atmospheric model, provided by M2020 EDL analysis team. Figure 7 shows the cumulative distribution functions for the eight candidate sites, with optimistic and conservative driving rates.

III. Analysis Results

A. Results Overview

Intuitively, CDFs in Figure 7 provide the probability of successfully reaching the ROIs within a specified driving distance and time. For example, on HOL with the optimistic driving speed assumption, there will be 80% chance of reaching the two ROIs within 10km of drive (Figure 7-(A)). We sought for 90% probability of meeting the BRS requirements (12 km, 85 Sols) with the conservative driving speed estimates. Figure 8 shows the 90th percentile distance and time, obtained from the CDFs. The ranges correspond to conservative and optimistic estimates. Only HOL and NIL violate the BRS requirements. Most notably, the conservative estimate of the 90th percentile driving

time for HOL is ~ 107 Sols, which is significantly greater than the BRS's 85 Sol requirement. As discussed in detail in Section III.B, this is mostly due to the sandy terrains.

Note that these results are specific to the placement of the ROIs and the landing ellipse of the Mars 2020 Project. Planners of future Mars rover missions should *not* take Figure 8 as the general traversability of each landing site. However, some traversability challenges listed on the rightmost column in Figure 8 are independent of ROI and ellipse placement, such as sandy surface on HOL, rock abundant terrain on JEZ, surface roughness on MAW, and scarps in SWM. These factors would likely be relevant to the traversability of future Mars rovers.

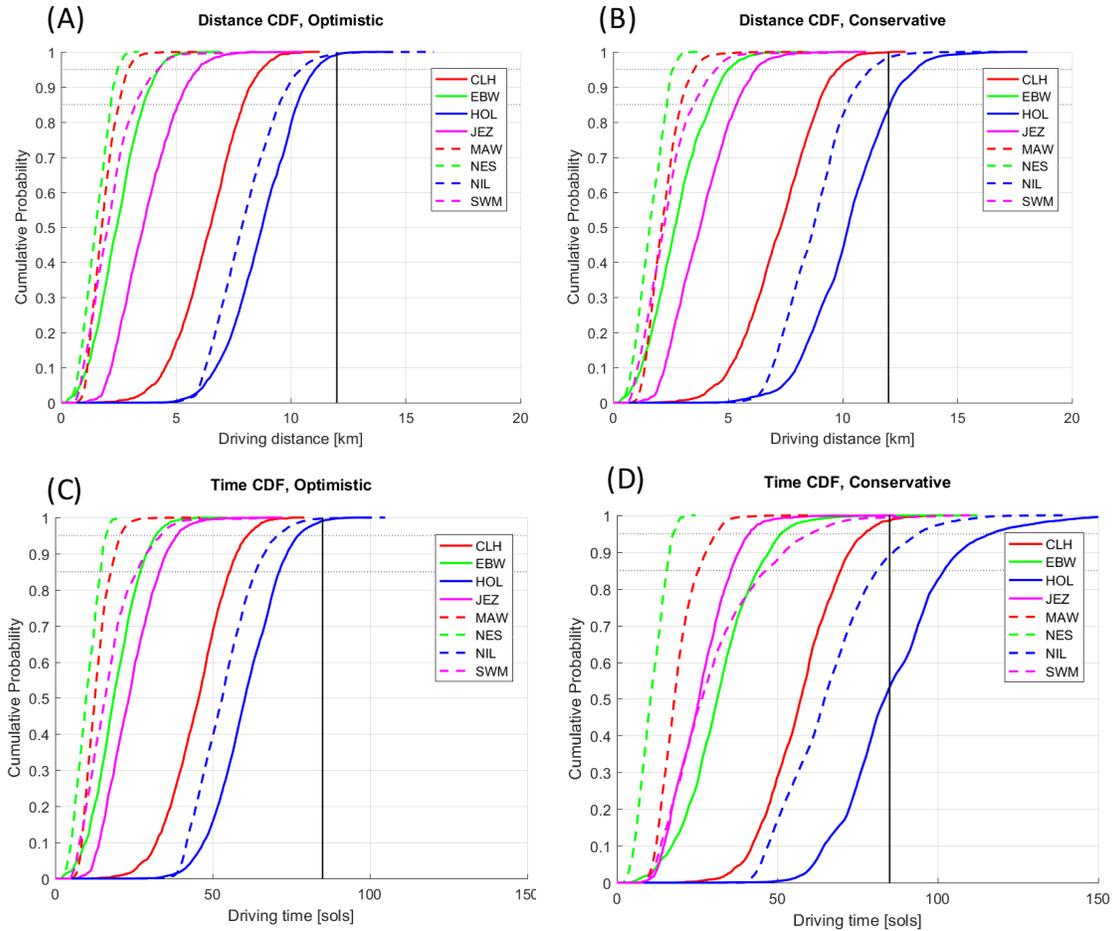


Figure 7 Cumulative distribution functions of (A) the driving distance with the optimistic driving speed estimates, (B) the driving distance with the conservative driving speed estimates, (C) the driving time with the optimistic driving speed estimates, and (D) the driving time with the conservative speed estimates.

	90% Time [Sol]	90% Distance [km]	Traversability challenges
BRS	85	12	(Baseline reference scenario)
CLH	57.7 – 72.7	8.3 – 9.3	Go-to site
EBW	28.9 – 47.6	3.8 – 4.6	Mantling unit with ripples Scarps on delta
HOL	73.7 - 106.8	10.6 – 12.5	Go-to site; >60% covered by potentially no-Autonav ripples; highways exist but in unfavorable directions Access to ROI (layered deposit) challenging due to high slope/sand
JEZ	35.5 – 38.1	5.5 – 5.8	High CFA on SE of ellipse but ROIs are on NW
MAW	19.1 – 28.0	2.7 – 3.2	Surface roughness could limit the speed of Autonav, but can achieve mission with conservative estimate
NES	15.1 – 16.5	2.3 – 2.4	Buttes and sand deposits, but localized and easy to go around
NIL	66.7 – 86.7	9.9 – 10.6	Go-to site Ripples but mitigated by highway in the favorable direction
SWM	29.6 – 52.5	3.7 - 4.0	Scarps, but traversable routes seem to exist across

Figure 8 Summary of the analysis results

B. Interpretation

Figure 9 shows the average driving distance (left) and time (right) and their breakdown by terrain type at each landing site. Among the eight landing sites, Holden turns out to be the worst both in terms of driving distance and time. With the conservative driving rate assumptions, Holden has about 85% and 55% chance to complete the inter-ROI drive with required distance and time, respectively. This is mainly because of two reasons. First, Holden is a “go-to” site, meaning that one of the ROIs is outside of the landing ellipse. The presence of landing hazards in and near the ROIs prevents the landing ellipse to be placed on the ROIs. The second reason is that, as we mentioned, the majority of the site is covered by slow-driving sandy ripples. In fact, as Figure 9 shows, about 27% of driving time is spent on sandy terrain (class 3 in Figure 5), even though it accounts for only ~6% of the distance.

Likewise, Nili Fossae and Columbia Hills have relatively long driving distance and time. They are also “go-to” sites. However, in case of Nili, the “highway” runs in a favorable direction and the ROIs of one type are widely distributed over the landing ellipse so that the rover can stop by the nearest one and take the highway to the ROI of the other type. Classes 1 and 2 (see Figure 5) constitute 98.5% of the distance and 92.9% of the time. In case of Columbia Hills, although it is categorized as a “go-to” site, the two ROIs are on the edge of the landing ellipse and closely situated to each other. Therefore, stopping by the first ROI requires almost no detour. Over 99% of the distance and time fall into classes 1 and 2. For these reasons, Nili Fossae and Clumbia Hills are not as worse as Holden.

The three middle-ranking sites are SW Melas, Eberswalde, and Jezero. The 90th percentile driving distance and time for these sites range approximately from 4-6 km and 30-50 Sols, respectively. At SW Melas and Eberswalde, the two ROIs are located in proximity of the center of landing ellipse while the terrain is mostly rough outcrop. As shown in Figure 9, with the conservative driving rate assumptions, 80% of the driving time at SW Melas and 63% of the driving time at Eberswalde are spent on rough terrains (class 2). SW Melas, located in Valles Marineris, has numerous scarps (Figure 10-left), but judging from visual inspection, traversable routes seem to exist to go across most of the scarps. Eberswalde has matling units covered by sand ripples (Figure 10-right). At Jezero, on the other hand, while ROIs are mildly off-centered towards the NW of the ellipse, benign terrains (class 1) constitute 97% of the distance and 91% of the time. For these reasons, although the driving distance is shorter at SW Melas

and Eberswalde, the driving time turns out to be shorter at Jezero, particularly in conservative cases where the penalty on rough terrains is high. Future improvement of the landing accuracy would make Jezero a more preferable landing site by mitigating the distance drawback.

Finally, Mawrth and NE Syrtis turn out to be the most favorable landing sites for surface traverse. In most cases, the driving distance and time would not exceed 3 km and 30 Sols, respectively. What these two sites have in common is that many ROIs are scattered over the ellipse so that any landing point could choose two (or three for MAW) nearby ROIs to visit, which significantly reduces the average driving distance. Figure 11 shows the average driving time breakdown by terrain at Mawrth (left) and NE Syrtis (right). Mawrth is covered mostly by rough outcrop, but the driving time is spent on diverse terrains: 61% on benign (class 1), 24% on rough (class 2), 14% on sandy (class 3). At NE Syrtis, the terrain is mostly smooth regolith, and so are the paths (96% of the distance and 92% of the time). Unlike Jezero, improving the landing accuracy is not likely to contribute to reducing the driving distance for these two sites because ROIs are evenly distributed over the ellipse.

As an example of the path-based CFA and slope statistics, Table 2 lists that of NE Syrtis. Summing up the 4 cells at the lower left corner, about 90% of the paths runs through CFA of less than 10% and slope of less than 10 degrees. In other words, 10% must get over CFA of >10% and/or slope of >10 degrees. This statistics is based on the optimistic driving rate estimates, meaning the statistics is shifted towards conservative because the routes are more likely to go through the complex terrains.

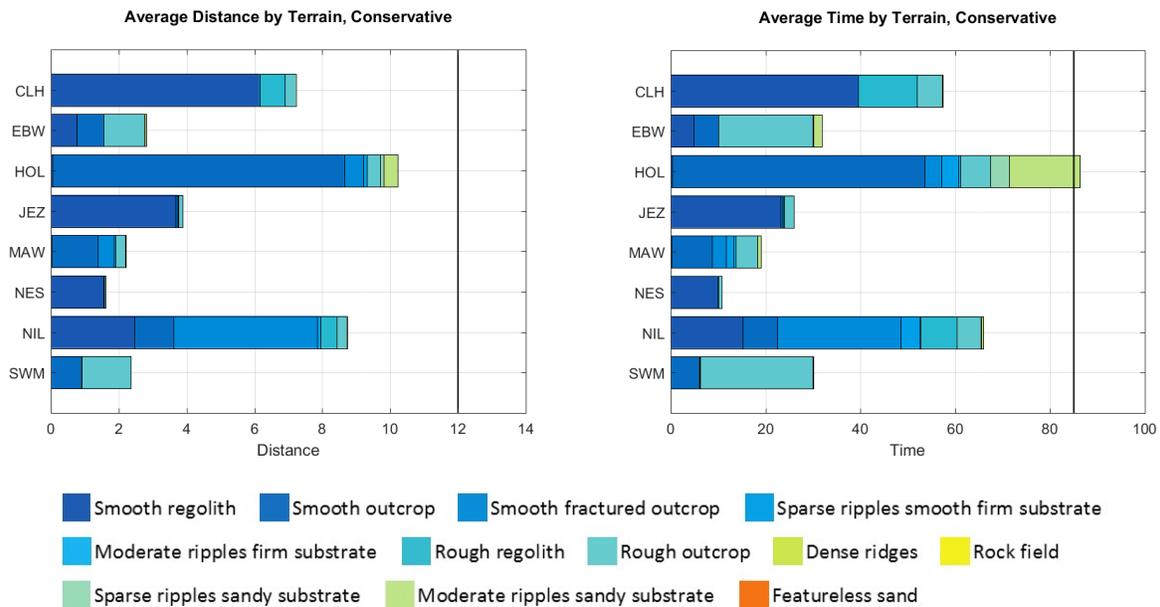


Figure 9. Average driving distance (left) and time (right) and their breakdown by terrain type at each landing site

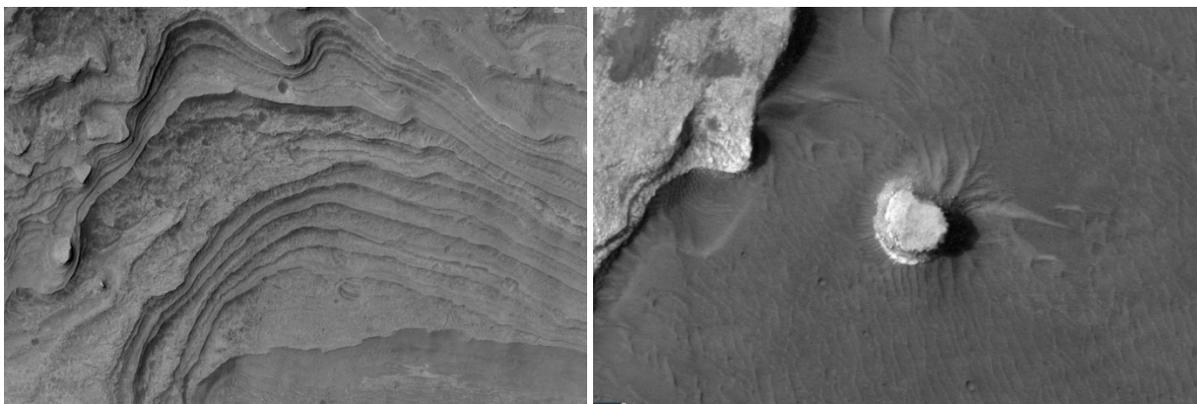


Figure 10 Scarps in SW Melas (left) and mantling unit covered by ripples in Eberswalde (right).

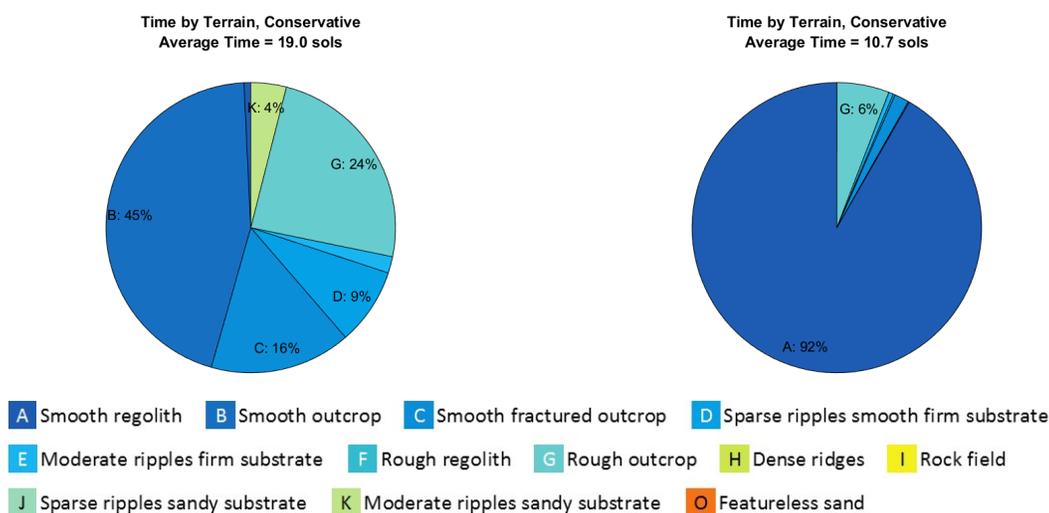


Figure 11 Average driving time breakdown by terrain type at Mawrth (left) and NE Syrtis (right).

Table 2 Path-based CFA and slope statistics at NE Syrtis

		CFA [%]			
		[0, 7)	[7, 10)	[10, 12)	[12, 15]
Slope [deg]	[15, 20]	0.42 %	0.60 %	0.07 %	0.01 %
	[10, 15)	4.50 %	1.63 %	0.17 %	0.06 %
	[5, 10)	22.85 %	3.40 %	0.33 %	0.69 %
	[0, 5)	58.01 %	6.16 %	0.52 %	0.56 %

C. Downselection of Landing Sites

The analysis approach presented in this paper was used to inform the public, the site proposers, and the Mars 2020 project science team about the expected mobility performance at each of the candidate landing sites. While the analysis produced a quantitative value for the number of sols required to drive between the Regions of Interest (ROIs), it is not expected that this is an actual prediction of mission performance in an absolute sense. In reality the mission will experience many twists and turns which will cause the mission timeline to deviate from the predicted timeline, as is expected for a mission of discovery. What this analysis does allow is a quantitative comparison of landing sites relative to each other. As the same analysis is applied uniformly to each landing site, it removes some of the human bias.

At the Mars 2020 Landing Site Workshop in February of 2017, the project downselected to three landing sites for further analysis; Columbia Hills, Jezero Crater, and Northeast Syrtis Major. For these sites, the ROIs will be refined with the site proposers and the project science team. The terrain classification, rock detection, and mobility model will also be refined as the team can focus on a smaller number of sites and the autonomous navigation capabilities of the rover are matured and better understood. This refined analysis will be used to aid in the selection of the final landing site in 2019. The tools will also be leveraged to aid in strategic route planning for the rover once it lands and will hopefully aid in orbital assessments of traversability for many Mars missions to come.

IV. Conclusion

This paper presented the analysis approach and results of the surface traversability analysis for the Mars 2020 Rover project. We developed the MTTT analysis toolset, consisting of various analytic capabilities including machine learning-based terrain classification, rock detection, and optimal route planning, to statistically evaluate the required traverse time and distance for visiting two ROIs. As a result of the analysis, we concluded that all the candidate landing sites except for two (NIL and HOL) satisfy the requirements on driving time and distance, even with the conservative estimate on rover's traverse speed. Among the remaining six sites, three sites (JEZ, NES, and CLH) were selected for further analysis in the Third Mars 2020 Landing Site Workshop, based on scientific values.

Acknowledgments

This work was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

References

- ¹Golombek, M. P., Cook, R. A., Moore, H. J. & Parker, T. J., 1997. Selection of the Mars Pathfinder landing site. *J. Geophys. Res. Plan.*, 102(E2), pp. 3967-3988.
- ²Golombek, M. P. et al., 2003. Selection of the Mars Exploration Rover landing sites. *J. Geophys. Res. Plan.*, 108(E12), p. 8072.
- ³Golombek, M. et al., 2012. Selection of the Mars Science Laboratory landing site. *Space Sci. Rev.*, 170(1-4), pp. 641-737.
- ⁴Golombek, M. et al., 2016. Selection of the InSight landing site. *Space Sci. Rev.*, pp. 1-91.
- ⁵Smith, D. E. et al., 2001. Mars Orbiter Laser Altimeter: Experiment summary after the first year of global mapping of Mars. *J. Geophys. Res. Plan.*, 106(E10), pp. 23689-23722.
- ⁶Neukum, G., Jaumann, R. & Team, H. C.-I. a. E., 2004. HRSC: The high resolution stereo camera of Mars Express. In: *Mars Express: the scientific payload*. Noordwijk: ESA Publications Division, pp. 17-35.
- ⁷Malin, M. C. et al., 2007. Context Camera Investigation on board the Mars Reconnaissance Orbiter. *J. Geophys. Res. Plan.*, Volume 112, p. E05S04.
- ⁸McEwen, A. S. et al., 2007. Mars Reconnaissance Orbiter's High Resolution Imaging Science Experiment (HiRISE). *J. Geophys. Res. Plan.*, 24(7), p. E05S02.
- ⁹Rothrock B, Papon J, Kennedy R, Ono M, Heverly M, Cunningham C., "SPOC: Deep Learning-based Terrain Classification for Mars Rover Missions," *AIAA SPACE 2016*, 2016.
- ¹⁰Long J, Shelhamer E, Darrell T., "Fully convolutional networks for semantic segmentation," *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015.
- ¹¹Chen LC, Papandreou G, Kokkinos I, Murphy K, Yuille AL., "DeepLab: Semantic Image Segmentation with Deep Convolutional Nets, Atrous Convolution, and Fully Connected CRFs," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017.
- ¹²Golombek, M. P., Huertas, A., Marlow, J., McGrane, B., Klein, C., Martinez, M., ... & Adams, D. (2008). Size - frequency distributions of rocks on the northern plains of Mars with special reference to Phoenix landing surfaces. *Journal of Geophysical Research: Planets*, 113(E3).
- ¹³Golombek, M., Huertas, A., Kipp, D., and Calef, F., 2012, "Detection and characterization of rocks and rock size-frequency distributions at the final four mars science laboratory landing sites," *Mars*, vol. 7, pp. 1-22.
- ¹⁴Golombek, M., and Rapp, D., 1997, "Size-frequency distributions of rocks on mars and earth analog sites: Implications for future landed missions," *Journal of Geophysical Research: Planets*, vol. 102, no. E2, pp. 4117-4129.
- ¹⁵Kirk, R. L. et al., 2008. Ultrahigh resolution topographic mapping of Mars with MRO HiRISE stereo images: Meter - scale slopes of candidate Phoenix landing sites. *J. Geophys. Res. Plan.*, Volume 113, p. E00A24.
- ¹⁶Burrough, P. A. & McDonell, R. A., 1998. *Principles of Geographical Information Systems*. New York: Oxford University Press.
- ¹⁷Ono, M., Rothrock, B., Almeida, E., Ansar, A., Otero, R., Huertas, A., and Heverly, M., 2016, "Data-Driven Surface Traversability Analysis for Mars 2020 Landing Site Selection," *Proc. IEEE Aerospace*