

Topographical Landmarks for Ground-Level Terrain Relative Navigation on Mars

Joshua Vander Hook, Russell Schwartz, Kamak Ebadi, Kyle Coble, Curtis Padgett

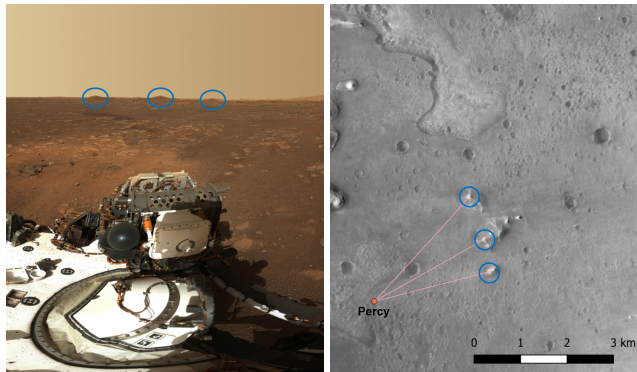
NASA Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Dr, Pasadena CA 91109

hook@jpl.nasa.gov, rschwa63@umd.edu, kamak.ebadi@jpl.nasa.gov,
kwc57@cornell.edu, curtis.w.padgett@jpl.nasa.gov

Abstract—Many of the tasks planned for future generation Mars rovers rely heavily on having accurate knowledge of the rover’s location in a Martian body-fixed coordinate system. Current solutions for localization require regular human intervention in order to detect and rectify drift, and thus stand to benefit from systems that can run in real-time on the rover itself. We study the feasibility and performance of an automated approach to localization in which the rover makes bearing-only measurements to geographic features in its surroundings (hills, boulders, peaked ridge-lines, etc.). When the location of these landmarks can be cross-referenced with a map of Mars, the resulting solution will be globally registered and will help correct any drift during visual-odometry-aided drives. This paper studies two related problems: First, how can we locate geographic features that the rover can feasibly see, when provided an elevation map of the surrounding terrain? We provide a software tool that can extract features from elevation maps for comparison to imagery. However, once these landmarks are identified, it is not obvious if a given path for the rover will contain sufficient features to navigate autonomously. Accuracy will depend on the quantity, range, and relative geometry of landmarks that are available. Thus, the second contribution is to provide a GIS plugin that analyzes the terrain informativeness of large operating areas as well as more specific paths. We present an analysis of Jezero Crater in which Perseverance’s Navcam is used as the hypothetical sensor. In certain favorable regions, worst-case localization accuracy in the 10 meter range is achieved from a single set of measurements (comparable to GPS on Earth). The map overlays generated by this analysis have the potential to aid in long-term mission planning by highlighting broad areas of high or low informativeness. These tools are computationally efficient and will be made open source to allow Mars mission planners, formulation studies, and rover drivers to plan for any future image-based self-localization capability.

TABLE OF CONTENTS

1. INTRODUCTION.....	1
2. RELATED WORKS	2
3. METHODOLOGY	3
4. RESULTS	3
5. TOOLS/PLUGINS	5
6. CONCLUSIONS.....	5
7. ACKNOWLEDGMENTS	5
REFERENCES	5
BIOGRAPHY	6



(a) on-board camera (b) digital elevation model

Figure 1: An image taken by Perseverance’s Mastcam-Z on sol 3 (cropped & stretched) showing three distinctly visible peaks and the corresponding topographic features marked on a map of Jezero Crater.

1. INTRODUCTION

For robotic tasks both on Earth and in space, having accurate knowledge of the robot’s location throughout the duration is essential in the completion of a wide range of missions. As a motivating example, the proposed Mars Sample-Return mission envisions the collection of a diverse set of samples from a variety of locations on the surface of Mars, in the hopes of later return to Earth. These criteria put substantial navigational demands on the Mars 2020 Perseverance Rover (theoretically responsible for the first leg of this mission) in terms of both range and speed [1]. For planetary rovers in general, since GPS is unavailable, localization can be achieved directly by matching surface features in the vicinity of the rover to surface imagery captured by onboard cameras. However, current implementations require regular human intervention in order to detect and rectify drift, and thus stand to benefit from systems that can run in real-time on the rover itself [2].

We consider an approach in which the rover makes a series of bearing-only measurements to known large-scale surface landmarks. Potential landmarks on the surface of Mars include large boulders in known locations (e.g., from HiRISE), the peaks of well-isolated hills, distant ridge-lines, and any other feature of the terrain that can be reliably identified both in Mars maps and located by the rover’s cameras. When the location of these landmarks can be cross-referenced with a map of the area, the resulting solution will be globally registered in the martian body-fixed coordinate system and will

help correct any drift during visual-odometry-aided drives. We seek to determine a rough upper bound on the effectiveness of this model of localization, depending on the terrain.

Under this model, the rover will be able to localize if there are at least two non-collinear landmarks available. In general, the accuracy of this localization depends on quantity, range, and relative geometry of landmarks that are available from the current location. A landmark is considered available when there is direct line-of-sight between the landmark and the camera sensor. Once a global set of landmarks is determined, we can associate a metric for terrain informativeness with each point on the map based upon the theoretical localization accuracy of a rover situated at that location. By computing this metric for every point in the scene, a raster is produced that captures the varying terrain informativeness for the entire region.

We present a complete pipeline for analyzing a digital elevation model (DEM) in order to assess the applicability of our landmark-based navigation model. First, potential landmarks are automatically detected in the terrain using standard morphological techniques. Then, the regions of visibility (“viewsheds”) of these landmarks are computed. Finally, the theoretical localization accuracy is determined at every point in the raster by computing the geometric dilution of precision (GDOP) in the resulting position estimate based on an assumed noise model in our primary sensor. This process is implemented as a set of plugins for the QGIS software package. We then present an analysis of Jezero Crater in which the Perseverance Rover’s Navcam is used as the hypothetical sensor. In certain favorable regions, worst-case localization accuracy in the 10s of meters is achieved (comparable to GPS on Earth). We show example map overlays generated by this analysis and discuss their potential to aid in long-term mission planning by highlighting broad areas of high or low informativeness.

An implementation of this approach would be relatively straightforward. It would include outfitting the rover with a static database of landmarks in its operational area, a system for detecting and recognizing landmarks within its field of view, and a general schedule/protocol for making the observations. The database might include the pre-computed viewsheds for each of the landmarks to aid in the recognition step. The resulting self-contained system would serve as an additional sensing modality which could help lower positional uncertainty during autonomous drives. An advantage of this approach is that the most computationally intensive steps - detecting landmarks and computing viewsheds - can be done entirely offline by computers on Earth.

Such an approach has not been studied extensively in this setting. Our main contributions can be summarized as follows:

- A novel method for procedurally extracting informative landmarks that uses both terrain morphology *and* viewshed analysis
- A well-defined metric for terrain informativeness (in the sense of information-gain from visible landmarks)
- A suite of open-source QGIS plugins for extracting the landmarks and computing the metric
- An analysis of Jezero Crater using these tools

2. RELATED WORKS

Terrain Relative Navigation

One well-studied method for localizing a robot (within a pre-defined region) is to compare observations to known terrain maps of the area. On Earth, this is often studied in the context of localizing an image taken in mountainous terrain [3], [4], [5]. These approaches use techniques such as skyline edge detection and contourlet matching. In [5], terrestrial pictures are aligned with topographic maps, using semantic segmentation of the query image. Typically these methods rely on a relatively-precise prior estimate of the GPS location where the image was taken and tend to be computationally demanding. A work highly relevant to our proposed application is [6], in which the authors present a method for building and maintaining a database of topographical landmarks that maximizes overall utility based on line-of-sight bearing measurements for use by planetary rovers. The authors use geometric dilution of precision (GDOP) as a metric for localization accuracy. Omitted from this work however is an explicit method for detecting these landmarks or for computing regions of visibility.

On Mars, terrain relative navigation was implemented successfully in the Mars 2020 Lander Vision System (LVS) which helped guide Perseverance to a safe landing sight during the descent stage [7]. Separate from the task of lander hazard avoidance, using terrain for rover localization is also well studied. In [8], various computer vision techniques are investigated for their applicability to terrain relative navigation at ground-level. Included in this work is a discussion of the stereo vision, visual odometry, and structure-from-motion systems used in the MER mission, which work by observing the rover’s immediate surroundings and allow for accurate local velocity estimation but are still subject to drift. The authors conclude that there is still work to do towards absolute localization relative to landmarks or orbital imagery. Other approaches have been explored such as [9], which considers aerial-to-ground localization using Mars helicopter aerial maps. Finally, [10] describes concurrent work at JPL that also features onboard processing of imagery for the purpose of localization. In this approach, semantic segmentation of captured imagery is compared to renders of the terrain in order to fix rover pose within the martian fixed-body coordinate frame.

Topographical Landmark Detection

The task of detecting recognizable topographical landmarks is a relatively old problem, with works like [11] presenting reliable methods for extracting peaks and ridges in terrestrial settings. However, the availability of high-accuracy topography data for Mars is relatively recent, coming in part from NASA’s Mars Reconnaissance Orbiter HiRISE telescope, and motivated largely by the need for high-accuracy DEM’s in the Mars 2020 LVS [12]. Much of the subsequent work to process these maps has been focused on either the LVS use-case [7] or on determining regions of likely wheel-slippage and overall traversability [13], [14].

Viewshed Analysis

Computing viewsheds is routine operation in geographic information systems and is used in a wide variety of applications. In [15], viewshed analysis is used to optimize the coverage of a network of radio masts. In [16], the author’s consider using viewsheds to inform future lunar exploration with an analysis of the errors introduced by uncertainty in the DEM, concluding that the boundary of the viewsheds can

be very sensitive to slight changes in elevation. Within the domain of robotics, viewshed analysis has been used in works like [17] to estimate the visibility of urban landmarks for use by autonomous vehicles. Pre-existing viewshed calculation tools are readily available and can be modified for use in our application.

3. METHODOLOGY

Our process for analyzing the informativeness of a section of terrain is comprised of three main steps:

1. Automatically detect potentially useful landmarks within the scene using terrain morphology.
2. Compute the regions of visibility (viewsheds) for these landmarks for a given observer height.
3. Using these viewsheds, for every point in the scene, determine which landmarks can be seen and their range to the rover. For a given measurement noise model, compute the theoretical localization accuracy (GDOP) at the point.

This result is a raster image in which each pixel represents the potential localization accuracy if the rover were to take measurements from the corresponding location. Alternatively, we can restrict our attention to a particular path within the scene representing a hypothetical drive campaign.

Landmark Detection

We extract salient landmarks from the DEM by the following process:

1. downsample DEM to ~ 50 m/px
2. classify terrain features with the `r.param.scale` utility from GRASS GIS
3. filter for regions identified as either `peak` or `ridge`
4. dilate and merge the resulting polygons (to reduce discretization artifacts)
5. extract point of maximal altitude for each polygon

This process generally results in landmarks that have the following desirable characteristics:

- approximately the right scale with respect to the range of motion of the rover
- widely visible throughout the scene (from ground level)
- readily identifiable within the rover’s camera frame

Empirically, each of these landmarks tend to correspond to a large boulder, a well-isolated hill, or the peak of a ridge-line. The set of candidate landmarks may be further refined in the next step.

Computing Viewsheds

In GIS literature, a viewshed is the set of points that are visible from a particular observer location via direct line-of-sight. Symmetrically, it is the set of points from which a particular landmark can be viewed [18]. We can use pre-existing viewshed calculation tools to produce a binary raster for each of our landmarks, in which the value of each pixel represents the visibility of the landmark from the corresponding location. These tools work on the principle of casting a ray from the landmark to every point in the scene, terminating when the ray first intersects the terrain, or when it reaches the horizon. We specifically use the open-source QGIS `Visibility Analysis` plugin, modified slightly to account for the elevation of the observer above the terrain when checking for ray intersections. In our case, we use the

height of perseverance’s mastcam (~ 3 m).

Additional considerations when producing these viewsheds include the curvature of the celestial body and any atmospheric refraction. Computing viewsheds for each landmark is by far the most computationally intensive step in our process. At this point, we may also want to discard unhelpful landmarks whose viewsheds are degenerate (in the sense that the intersection with the area of interest is too small to be of use) or redundant (in the sense that the viewshed is fully contained within the viewshed of another landmark, with no significant baseline between them).

Localization Accuracy

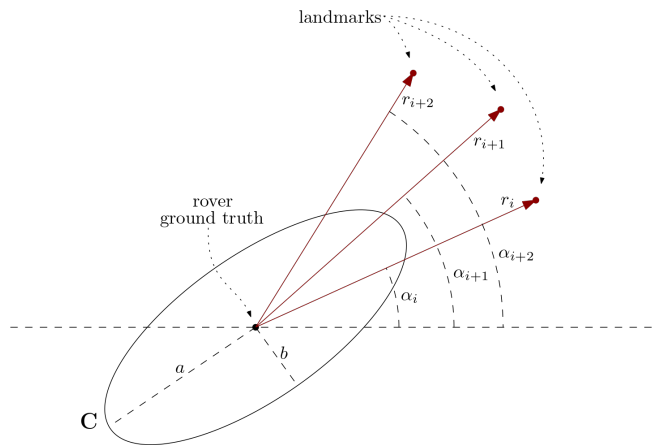
Finally, we iterate over every point in the scene and examine the corresponding points in the viewshed rasters in order to determine what subset of n landmarks will be visible. We then compute the range r_i and azimuth a_i from the rover to each landmark. We model our measurements as being draw from the distribution

$$\hat{a}_i = a_i + \mathcal{N}(0, \sigma_i^2)$$

which corresponds to an unbiased gaussian noise model with parameter σ_i , which we refer to as the pointing accuracy. The Cramer-Rao lower bound yields a covariance ellipse \mathbf{C} describing the distribution of resulting position estimates, which is given by

$$\mathbf{C}^{-1} = \sum_{i=1}^n \frac{1}{r_i \sigma_i} \begin{bmatrix} \sin^2 \alpha_i & -\sin \alpha_i \cos \alpha_i \\ -\sin \alpha_i \cos \alpha_i & \cos^2 \alpha_i \end{bmatrix}$$

Figure 2: The covariance ellipse \mathbf{C} resulting from measurements to landmarks with range r_i and azimuth a_i . Here, $\text{GDOP} = a + b$.



We take the average of the major and minor axis of as a quantification for size of this ellipse, which provides a reasonable metric for the accuracy of this position estimate. This corresponds with the standard notion of GDOP [19].

4. RESULTS

We ran our landmark-detection tool on a 30km (6000px) square elevation map of Jezero Crater (indeed, the same DEM

used in the Mars 2020 LVS). In total, 54 landmarks were automatically detected, primarily composed of small hills in the crater’s interior and sharp peaks in the ridgeline of the crater rim.

Table 1: Geographic distribution of landmarks detected in Jezero Crater DEM

Location	Landmarks Detected
Crater Interior	21
Crater Rim	13
Crater Exterior	20
Total	54




Figure 3: Topographic landmarks automatically extracted from a DEM of Jezero Crater (30km x 30km)

In order to gauge the pointing accuracy of bearing measurements made to these landmarks, we consider the following sources of error:

- the finite angular resolution of the camera
- the accuracy of the camera distortion model/reported pose of the instrument
- our ability to recognize landmarks within an image

We take one of the monocular Navcams in Perseverance’s main mast as our hypothetical sensor, which has a horizontal resolution of 0.33 mrad/px [20], and assume no appreciable noise in the distortion model or reported instrument pose. By performing a study of a terrestrial analog image set, we determined that landmarks can be reliably identified in an image to within a standard deviation of ~10px using standard computer vision techniques that can easily run onboard. In this study, we assumed the rover has at least a rough prior estimate of its location. By consulting the pre-computed viewsheds at this point, the rover prioritizes which landmarks it should look for in its field of view. Overall these values

yields a pointing accuracy of $\sigma \approx 3.3$ mrad, which we take to be fixed for all measurements.

We then ran our terrain informativeness plugin on the DEM using the automatically detected landmarks, the estimated pointing accuracy, and the proper value for the instrument height. The resulting raster shows a wide variety in accuracy, ranging from GDOP ≈ 10 m in the most informative regions (e.g. immediately to the northwest of the three hills noted in Figure 1) to GDOP > 100 m in the least informative regions (e.g. in the riverbed shown in Figure 5).

We also examine a hypothetical drive campaign by selecting a random path through the scene, starting at Perseverance’s landing site. As seen in Figure 6, the accuracy may change rapidly along a continuous path as certain landmarks become occluded or as new ones become visible.




Figure 4: Map overlay of Jezero Crater showing predicted localization accuracy (GDOP) at every point






Figure 5: In the riverbed, the uncertainty effectively becomes infinite as all landmarks become occluded by the banks on either side

Figure 6: Analysis of a hypothetical drive campaign through Jezero Crater



(a) A potential rover path with bearing measurements and resulting covariance ellipse at a particular location (40σ shown for clarity at large scale)



(b) The predicted localization accuracy throughout the campaign (capped at 85m) with the highest values occurring in the riverbed

5. TOOLS/PLUGINS

These tools are implemented as extensions for the open-source QGIS software package. A separate processing plugin exists for each main step: detecting landmarks, computing viewsheds, and creating the final raster. The user may specify the details of their application, including various aspects of the desired landmarks, observer height, surface curvature, atmospheric refraction, pointing accuracy, and estimator quality metric. An additional plugin was developed that accepts an arbitrary path through the scene (representing a hypothetical drive campaign) and produces a video animation of the rover moving along this path, highlighting the visible landmarks at each point and rendering the resulting covariance ellipse. For access to the suite of plugins, contact the authors at hook@jpl.nasa.gov.

6. CONCLUSIONS

We consider a method for rover localization using topographical landmarks and establish a rough upper bound on estimator quality. Our results are mixed but promising. There are significant regions of the martian surface where

our estimator achieves GDOP values in the 10m range from only a single set of observations. However, there are also some large areas and numerous small pockets where quality drops off significantly due to a lack of visibility. This is an inevitable feature of landmark-based localization, and renders our approach unsuitable as a sole method of local navigation. The utility lies in the potential to augment existing techniques by providing sporadic position estimates that are globally registered, helping to correct any accumulated drift.

Our approach demonstrates other advantages. For instance, since expected estimator quality can be pre-computed, mission planners and formulation studies can generate maps like Figure 4 that highlight favorable and unfavorable regions in order to inform their choice of route. Additionally, the computational demands involved in visually detecting landmarks and producing a solution are moderate. Indeed the Mars 2020 LVS demonstrated that tasks of this nature can be feasibly performed on-board.

7. ACKNOWLEDGMENTS

The research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004). Thanks to the Mars Perseverance entry, descent, and landing team for furnishing Mars DEMs used in this work. Particular thanks to Andrew Johnson.

REFERENCES

- [1] B. K. Muirhead, A. Nicholas, and J. Umland, “Mars sample return mission concept status,” in *2020 IEEE Aerospace Conference*, 2020, pp. 1–8.
- [2] C. Wong, E. Yang, X.-T. Yan, and D. Gu, “Adaptive and intelligent navigation of autonomous planetary rovers — a survey,” in *2017 NASA/ESA Conference on Adaptive Hardware and Systems (AHS)*, 2017, pp. 237–244.
- [3] G. Baatz, O. Saurer, K. Köser, and M. Pollefeys, “Large scale visual geo-localization of images in mountainous terrain,” in *Computer Vision – ECCV 2012*, A. Fitzgibbon, S. Lazebnik, P. Perona, Y. Sato, and C. Schmid, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 517–530.
- [4] L. Baboud, M. Čadík, E. Eisemann, and H.-P. Seidel, “Automatic photo-to-terrain alignment for the annotation of mountain pictures,” in *CVPR 2011*, 2011, pp. 41–48.
- [5] G. Baatz, O. Saurer, K. Köser, and M. Pollefeys, “Leveraging topographic maps for image to terrain alignment,” in *2012 Second International Conference on 3D Imaging, Modeling, Processing, Visualization Transmission*, 2012, pp. 487–492.
- [6] T. J. Steiner, T. M. Brady, and J. A. Hoffman, “Graph-based terrain relative navigation with optimal landmark database selection,” in *2015 IEEE Aerospace Conference*, 2015, pp. 1–12.
- [7] A. Johnson, S. Aaron, J. Chang, Y. Cheng, J. Montgomery, S. Mohan, S. Schroeder, B. Tweddle, N. Trawny, and J. X. Zheng, “The lander vision system for mars 2020 entry descent and landing,” 2017.
- [8] L. Matthies, M. Maimone, A. Johnson, Y. Cheng, R. Willson, C. Villalpando, S. Goldberg, A. Huertas,

A. Stein, and A. Angelova, "Computer vision on mars," *International Journal of Computer Vision*, vol. 75, pp. 67–92, 10 2007.

- [9] K. Ebadi and A.-A. Agha-Mohammadi, "Rover localization in mars helicopter aerial maps: Experimental results in a mars-analogue environment," in *Proceedings of the 2018 International Symposium on Experimental Robotics*, J. Xiao, T. Kröger, and O. Khatib, Eds. Cham: Springer International Publishing, 2020, pp. 72–84.
- [10] K. Ebadi, K. Coble, D. Kogan, D. Atha, R. Schwartz, C. Padgett, and J. Vander Hook, "Semantic mapping in unstructured environments: Toward autonomous localization of planetary robotic explorers," in *2022 IEEE Aerospace Conference*, 2022.
- [11] I.-S. Kweon and T. Kanade, "Extracting topographic features for outdoor mobile robots," *Proceedings. 1991 IEEE International Conference on Robotics and Automation*, pp. 1992–1997 vol.3, 1991.
- [12] R. L. Kirk, D. P. Mayer, R. L. Ferguson, B. L. Redding, D. M. Galuszka, T. M. Hare, and K. Gwinner, "Evaluating stereo digital terrain model quality at mars rover landing sites with hrsc, ctx, and hirise images," *Remote Sensing*, vol. 13, no. 17, 2021. [Online]. Available: <https://www.mdpi.com/2072-4292/13/17/3511>
- [13] G. Hedrick and Y. Gu, "Terrain-aware traverse planning for a mars sample return rover," *Advanced Robotics*, vol. 0, no. 0, pp. 1–16, 2021. [Online]. Available: <https://doi.org/10.1080/01691864.2021.1955000>
- [14] M. Ono, B. Rothrock, E. Almeida, A. Ansar, R. Otero, A. Huertas, and M. Heverly, "Data-driven surface traversability analysis for mars 2020 landing site selection," in *2016 IEEE Aerospace Conference*, 2016, pp. 1–12.
- [15] Y.-H. Kim, S. Rana, and S. Wise, "Exploring multiple viewshed analysis using terrain features and optimisation techniques," *Computers Geosciences*, vol. 30, no. 9, pp. 1019–1032, 2004. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0098300404001372>
- [16] P. Mahanti and M. S. Robinson, "Informed lunar exploration with probabilistic viewsheds," in *European Planetary Science Congress*, vol. 9, Apr. 2014, pp. EPSC2014–620.
- [17] M. Aleksandrov, S. Zlatanova, L. Kimmel, J. Barton, and B. Gorte, "Voxel-based visibility analysis for safety assessment of urban environments," *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. IV-4/W8, pp. 11–17, 2019. [Online]. Available: <https://www.isprs-ann-photogramm-remote-sens-spatial-inf-sci.net/IV-4-W8/11/2019/>
- [18] M. Nachon, S. Borges, R. C. Ewing, F. Rivera-Hernández, N. Stein, and J. K. Van Beek, "Coupling mars ground and orbital views: Generate viewsheds of mastcam images from the curiosity rover, using arcgis@ and public datasets," *Earth and Space Science*, vol. 7, no. 9, p. e2020EA001247.
- [19] R. B. Langley, "Dilution of precision," *GPS World*, vol. 10, no. 5, pp. 52–59, 1999.
- [20] J. Maki, D. Gruel, C. McKinney, M. Ravine, M. Morales, D. Lee, R. Willson, D. Copley-Woods,

M. Valvo, T. Goodsall, J. McGuire, R. Sellar, J. Schaffner, M. Caplinger, J. Shamah, A. Johnson, H. Ansari, K. Singh, T. Litwin, and S. Algermissen, "The mars 2020 engineering cameras and microphone on the perseverance rover: A next-generation imaging system for mars exploration," *Space Science Reviews*, vol. 216, 12 2020.

BIOGRAPHY



Dr. Joshua Vander Hook received a PhD in Computer Science at the University of Minnesota in 2015. His research involved designing autonomous robots that can assist in surveying and data-gathering in remote areas. He held a Doctoral Dissertation Fellowship at the University of Minnesota.



Russell Schwartz is an undergraduate at the University of Maryland in his senior year, scheduled to graduate with a dual B.S. in Mathematics and Computer Science in Spring 2022. He is also an intern at JPL and plans on pursuing a Master's degree in Robotics beginning the following Fall.



Dr. Kamak Ebadi is currently a Robotics Technologist at NASA JPL, Pasadena, CA, USA. He received his Ph.D degree in Computer and Electrical Engineering from Santa Clara University in 2020, and was a Postdoctoral Fellow with NASA JPL-California Institute of Technology from 2020 to 2021. His research interests include computer vision, multi-robot perception and autonomy in perceptually degraded and extreme environments.



Kyle Coble received his B.S. degree in Mechanical Engineering in 2016 from Cornell University. He is scheduled to complete an M.S. degree in Systems, Control, and Robotics from the KTH Royal Institute of Technology in late 2021. His interests include computer vision and state estimation for autonomous robots in extreme environments.



Dr. Curtis Padgett is the Supervisor for the Maritime and Aerial Perception Systems Group. He completed his doctoral work in pattern recognition while working at JPL. His interests include computer vision, structure from motion, automated calibration, pattern classification, star identification and machine learning.