

# Censible: A Robust and Practical Global Localization Framework for Planetary Surface Missions

Jeremy Nash<sup>1</sup>, Quintin Dwight<sup>1</sup>, Lucas Saldyt<sup>1</sup>, Haoda Wang<sup>1</sup>, Steven Myint<sup>1</sup>, Adnan Ansar<sup>1</sup>, Vandi Verma<sup>1</sup>  
<sup>1</sup> Jet Propulsion Laboratory, California Institute of Technology, USA

**Abstract**—To achieve longer driving distances, planetary robotics missions require accurate localization to counteract position uncertainty. Freedom and precision in driving allows scientists to reach and study sites of interest. Typically, rover global localization has been performed manually by humans, which is accurate but time-consuming as data is relayed between planets. This paper describes a global localization algorithm that is run onboard the Perseverance Mars rover. Our approach matches rover images to orbital maps using a modified census transform to achieve sub-meter accurate, near-human localization performance on a real dataset of 264 Mars rover panoramas. The proposed solution has also been successfully executed on the Perseverance Mars Rover, demonstrating the practicality of our approach.

## I. INTRODUCTION

The NASA Perseverance rover landed on Mars on February 18, 2021. It is part of the planned Mars Sample Return campaign to collect sample tubes and return them to Earth for scientific analysis. Collecting a diverse set of sample tubes requires traversing long distances across the Martian surface.

It is the first rover to be able to drive autonomously at close to its maximum electromechanical speed and its self-driving autonomous navigation system has been used to evaluate 88% of the distance traveled [1]. The greatest distance it has driven without human review is 699.9m over three days which is a planetary rover record. This overall distance is limited by growth in position uncertainty, shown in Fig. 1.

The position uncertainty is reset by localizing the rover on a global map of Mars. To estimate absolute orientation, the rover periodically runs a sun finding activity. However, to reset the rover's position uncertainty, human rover operators downlink panorama images and match them to an orbital map manually. Automating position estimation would enable the localization process to be fully performed on Mars removing any limitation on drive distance due to uncertainty growth.

This is the first paper to solve this problem on real Mars data with sub-meter accuracy and no outliers, while also being compatible with the typical limitations of a flight mission like limited compute, memory, uplink data volume, as well as thermal and power constraints. This paper advances the state-of-practice for planetary robots with:

- A novel model-based approach that achieves state-of-the-art, near-human localization performance.

Part of this research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. © 2023. All rights reserved.

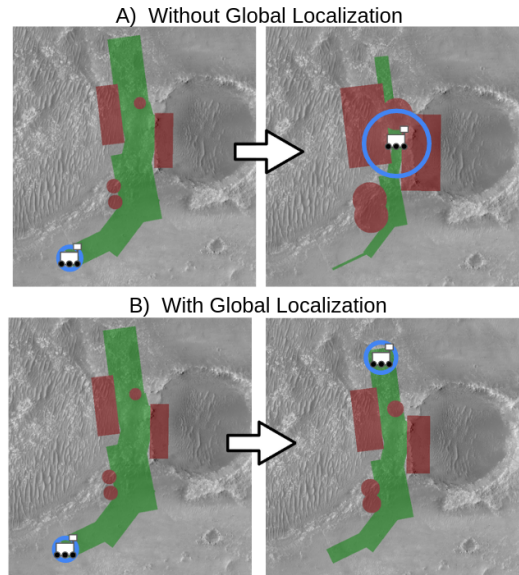


Fig. 1: A) On Sol 385, the Mars Perseverance Rover attempted a long drive through a narrow corridor, but the growth in its odometry position uncertainty (blue) resulted in the dilation of the hazard map (red) and an early termination of the drive. B) Our framework, Censible, addresses this common problem through accurate and reliable localizations to an orbital reference map, allowing the rover to reduce position uncertainty, shrink the hazard map, and find a path.

- Comprehensive experiments on real Mars rover data, and a demonstration on the Perseverance Mars Rover showing the practicality of our approach.

## II. RELATED WORK

**In current practice** on all Mars rover missions, global localization is performed manually by humans. The rover drives until its odometry solution accumulates too much uncertainty and captures a 360 panorama which human experts will manually match to an orbital map [2].

**Image matching approaches** cast global localization as an image registration problem between the rover images and the orbital map. Previous image matching approaches include normalized cross-correlation [3], [4], [5], sum-of-squared differences, and mutual information [6], [7]. These approaches generally have a significant outlier rate, preventing adoption in real missions. Particle filters have been proposed to address this, but require images across multiple locations [4].

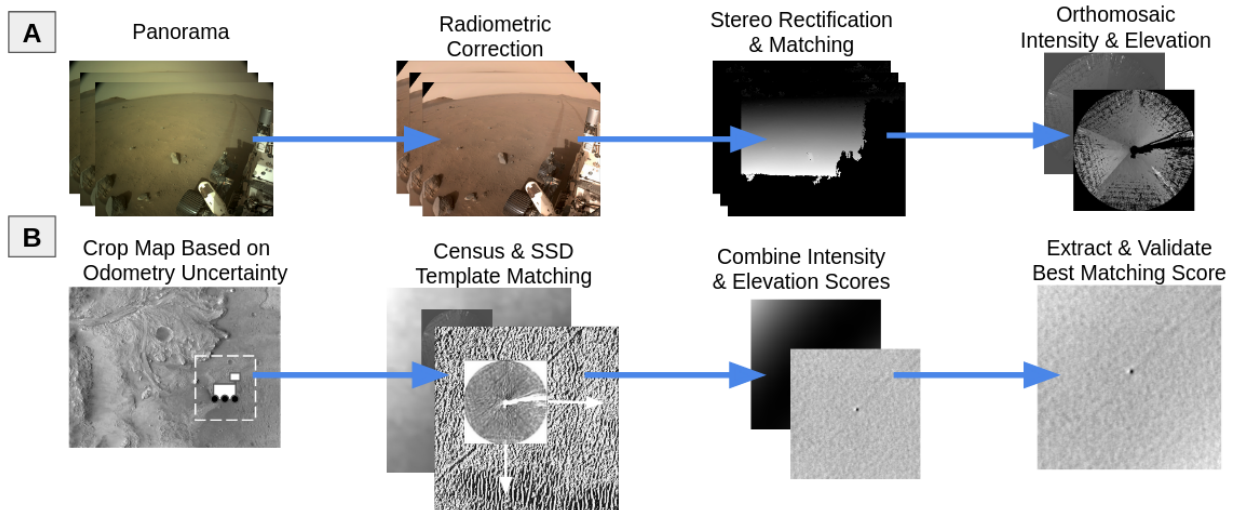


Fig. 2: High-level overview of the global localization framework. A) Transforming the rover images to match the orbital map: a 360 rover panorama of stereo images are radiometrically corrected, stereo matched, and orthographically projected into a 2D grid. B) Matching the rover images to the orbital map: the orbital map is cropped based on rover odometry’s position uncertainty, the orthomosaics are template matched to the orbital map based on the census and SSD similarity measures, and the best score is extracted and converted into a delta position update.

**Feature-based approaches** extract and match image features between rover and orbital images. Common approaches include extracting SIFT features [8], terrain-specific features like craters [9] or rocks [10], or some combination of the two [11]. These approaches naturally require terrain with unique features, and they struggle in low feature terrain that the rovers often drive through. A related category uses bundle adjustment to match features across many rover images to reduce odometry error [12], [13], although the error still fundamentally grows with distance traveled.

**Horizon-based approaches** match the local horizon to surface digital elevation models (DEMs), and various methods have been proposed to extract the horizons [14], [15], [16], [17], [18] or instead match mountain peaks [19], [20]. Accuracy is typically on the order of 100m, at best 10m. These methods are valuable for orientation estimation, but their accuracy does not match the position accuracy requirements of a typical Mars rover mission scenario.

**Deep CNN approaches** are a compelling new category that train a deep network to match rover images to orbital images [21], [22], [23]. Many of these rely on significant amounts of training data, often relying on simulation for both training and evaluation. The large size of the networks also presents a practical challenge for integration into a real rover mission, both in terms of Earth-Mars uplink volume and rover memory and computation constraints.

### III. METHODOLOGY

Our *orbital maps* are based on HiRISE, which is the highest resolution camera orbiting Mars with 0.25m image resolution for the red/NIR channel. The DEMs are derived from stereo matching and have 1.0m horizontal resolution. The Perseverance rover has wide-angle lens color stereo cameras mounted on a pan-tilt mast for navigation. The pan-



Fig. 3: A rover image and the orbital map rendered from the same perspective. Even with the highest resolution orbital maps available, there are significant differences in appearance, geometry, and shadows between them. Our approach produces robust matches despite these differences.

tilt can capture a 360 panorama of stereo images around the rover, which forms the input for our method.

We model rover localization as an image registration problem between the rover images and the orbital map. We take a model-based approach to make the rover images match the orbital map image as closely as possible, then find a similarity measure that is invariant to the differences. The following subsections detail the stages of our approach, depicted in Fig. 2.

**Absolute Orientation Estimation:** Rover missions can typically estimate absolute orientation with high accuracy, either by estimating rover orientation based on the position of the sun, the stars, or internally through gyrocompassing. By running one of these activities in advance, the registration problem can be simplified to a search over 2D translation.

**Panoramic Stereo Imaging:** We acquire a 360 panorama of stereo pairs around the rover to maximize the area of imaged terrain. For Perseverance, we use monochrome red channel images, which are spectrally the most similar to the red/near-infrared channel HiRISE map.




Fig. 4: The locations of all 264 panoramas localized in Jezero Crater on Mars, represented by white dots. The panorama dataset captures the entire rover mission to date with a wide variety of terrains over 2.5 years. This includes featureless terrain, complex terrain with occlusions, terrain modified due to wheel tracks, and a variety of shadows and lighting conditions.




Fig. 5: Comparison of sun angles for the rover images (red) versus the orbital map images (blue). The solid black lines represent the sun angle boundaries in Jezero Crater (summer and winter solstice), and the dotted black line is the spring/fall equinox.

**Radiometric Correction:** We model the image formation process from pixel values to scene radiances. A key motivation for this stage is to make the lighting consistent across all images in the combined orthomosaic. Common sources of lighting discrepancies include vignetting from the rover lens, or differences in image exposures.

$$p(x, y) = \tau_{comp}(f(x, y)r_s(x, y)t_{exp}) + b \quad (1)$$

represents the function from scene radiances  $r_s$  to pixel values  $p(x, y)$ , where  $\tau_{comp}$  is a square-root companding table to compress 12 bits to 8 bits,  $t_{exp}$  is the exposure time,  $b$  is a camera-specific bias, and  $f(x, y)$  is the flat field.

We model the flat field  $f(x, y)$  as a deformable radial polynomial [24]. This step removes a significant vignette at the edges of the wide-angle stereo images which is a nuisance for image matching. We use this parametric model instead of a raw flat-field image to minimize uplink data volume to the rover, on the order of tens of megabytes:

$$r_c(x, y) = (x - x_0)^2 + (\eta(y - y_0))^2$$

$$f(x, y) = \sum_{n=0}^3 c_n r_c(x, y)^n \quad (2)$$

where  $r_c(x, y)$  is the radius from column and row offsets  $x_0$  and  $y_0$  (roughly the image center),  $\eta$  is the non-radiality term which allows the curve to be asymmetric, and  $c_n$  are the coefficients of a third-degree polynomial. These parameters are estimated using nonlinear least squares minimization against the raw flat field correction images.

The image formation model above is inverted to obtain the scene radiances, where  $\tau_{comp}^{-1}$  is the decompanding table, the inverted lookup-table for  $\tau_{comp}$ :

$$r_s(x, y) = \frac{\tau_{comp}^{-1}(p(x, y)) - b}{f(x, y)t_{exp}} \quad (3)$$

This model notably does not capture lens flare, which is a nuisance present in images pointed towards the sun, but generally difficult to model and remove. Since panorama images are captured within ten minutes of each other, variations in dust optical depth and sun elevation did not have a significant enough impact to warrant inclusion in this model.

**Stereo Matching:** Lens distortion is modeled with a temperature-dependent CAHVORE model, a fisheye lens distortion model [25]. Images are stereo rectified, and stereo matching is performed using semi-global matching [26]. The SGM model implicitly assumes smoothly varying terrain geometry, which is a reasonable assumption for Martian terrain. SGM is particularly helpful in scenes with low texture, either due to inherently featureless terrain or poor lighting conditions due to the opposition effect (also known as shadow hiding). The algorithm uses a Sobel pre-filter, 5x5 blocks with a SAD-based correlator, and matching is performed with the full-scale two-pass algorithm in 8 directions. Quadratic subpixel estimation is used for more accurate ranges, and a series of post-filters removes erroneous matches: a pseudo left-right line-of-sight consistency check, a uniqueness ratio check, and a speckle post-filter.

**Orthoprojection:** The stereo point cloud is rotated to align with the orbital frame, using the previously recovered rover absolute orientation and the mast joint angles. These points are filtered to between 3 and 40 meters to remove the rover body and shadows in the near range, and noisy stereo in the far range. The stereo point clouds are ortho-projected into a 25cm resolution pixel grid to match the orbital appearance maps. The 1m resolution DEM is interpolated to a 25cm resolution grid to match the appearance map resolution to facilitate adding the scores together later. Pixels that fall into same grid cell are averaged.

**Template Matching:** To establish the rover's position, the rover orthomosaic image is matched against the orbital map image. The orbital map is cropped around on the rover's position estimate from odometry, with a width of the rover's position uncertainty plus the stereo orthomosaic width. Both the template and orbital map undergo a modified census transform, a non-parametric local image transform that is invariant to monotonic variations in intensity. The transform's partial invariance to brightness differences between wedges is also particularly useful for any unmodeled radiometric effects like lens flare.

The census transform is defined by a comparison operator between nearby pixel intensities. The proposed implementation accounts for missing data by adding a third category, which is ignored by the modified Hamming distance (Fig. 6).

$$c(p, p') = \begin{cases} 0 & \text{if } p > p' \\ 1 & \text{if } p \leq p' \\ \emptyset & \text{if either } p \text{ or } p' \text{ have no data} \end{cases} \quad (4)$$

This comparison is performed over a local window of pixels  $W$ , which is typically  $3 \times 3$ . Self-comparisons are excluded, resulting in 8 comparisons:

$$W = \{(i, j) \mid i, j \in \{-1, 0, 1\}, (i, j) \neq (0, 0)\} \quad (5)$$

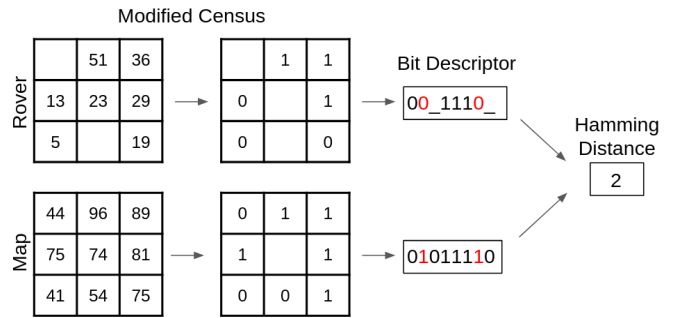


Fig. 6: Modified census transform with missing data, which is not penalized by the modified Hamming distance.

A typical census transform results in an 8-field descriptor for each pixel location  $(x, y)$  in an image  $I$  [27]. This transform is performed on both the rover orthomosaic and orbital map, resulting in a descriptor  $D$  for each pixel location:

$$D(x, y) = c(I(x, y), I(x + i, y + j)) \quad i, j \in W \quad (6)$$

The proposed algorithm uses a modified Hamming distance which does not penalize missing data, which comes from geometric occlusions and stereo failures that are a natural consequence of the rover's perspective:

$$H(d, d') = \begin{cases} 0 & \text{if } d = d', d = \emptyset \text{ or } d' = \emptyset \\ 1 & \text{otherwise } (d \neq d') \end{cases} \quad (7)$$

The Hamming distance between two descriptors sums over each comparison within the window  $W$ :

$$H(D, D') = \sum_{k=1}^{|W|} H(d, d') \quad (8)$$

The census score matrix  $S_C$  is computed by sliding the rover descriptor image  $D_r$  over the map descriptor image  $D_m$  and summing the modified Hamming distances:

$$S_C(x, y) = \sum_{i, j} H(\underbrace{D_m(x + i, y + j)}_{map}, \underbrace{D_r(i, j)}_{rover}) \quad (9)$$

The elevation score matrix uses a sum of squared distances (SSD) similarity measure between the rover  $E_r$  and map  $E_m$  elevations. For invariance to absolute elevation differences, the average elevation is subtracted from each window:

$$S_E(x, y) = \sum_{i, j} \left( \underbrace{(E_m(x + i, y + j) - \mu_{m_{x,y}})}_{map} - \underbrace{(E_r(i, j) - \mu_r)}_{rover} \right)^2 \quad (10)$$

The objective of template matching is to establish which  $x, y$  location minimizes  $S_C$  and  $S_E$ , indicating that the template and image region match. This search occurs over discrete pixels that represent 25cm cells. A typical 30m search range results in a  $240 \times 240$  pixel search region:

$$\min_{x, y \in 1..240} S_C(x, y) + \alpha S_E(x, y) \quad (11)$$

where  $\alpha$  enables a weighted sum of the census and elevation scores, with greater weight placed on the census scores. The minimizing  $(x, y)$  pixel coordinates are then converted to meters and then a global northing and easting.

**Subpixel Estimation:** The accuracy of the approach above is fundamentally limited by map resolution, so we estimate the subpixel location by fitting a 2D quadratic surface model to the 3x3 pixel neighborhood surrounding the best score:

$$p(x, y) = \sum_{i=0}^2 \sum_{j=0}^2 a_{i,j} x^i y^j \quad (12)$$

The quadratic model parameters are estimated using least-squares, and the subpixel location is recovered at the minimum critical point. In the rare case that the quadratic model is a poor fit to the data, a 1D quadratic model is fit to the 3x1 row and 1x3 column neighborhood, and the subpixel row and column are estimated independently. The minimum principal curvature of the 2D quadratic surface is additionally calculated as a confidence measure, based on an eigendecomposition of the 2D quadratic Hessian matrix.

**Map Distortion Model:** All map projections from sphere to plane have an inherent but modelable distortion. The equidistant cylindrical map projection has an inherent scaling error in the east-west direction that grows the farther the rover travels away from the map’s latitude of true scale  $\varphi_{ts}$ . For Perseverance, this could amount to a 1m discrepancy between map and rover every 500m. To correctly translate map vectors to rover vectors, the east-west scaling error can be compensated by the formula:

$$eastings_{map} = \frac{eastings_{rover} * \cos \varphi_{ts}}{\cos \varphi} \quad (13)$$

where  $\varphi$  is the rover’s latitude and  $\varphi_{ts}$  is the latitude of true scale (or standard parallel). No scale correction is required for the northing component.




Fig. 7: Localization accuracy of our approach on the benchmark dataset, using a 30m radius search window (the max expected rover uncertainty). Our approach achieves near-human performance with no significant outliers.

## IV. EXPERIMENTS

### A. Datasets

Our benchmark dataset consists of all available navigation camera panoramas from the Perseverance Mars Rover mission, 264 in total. The dataset captures a wide variety of environments from the past 2.5 years of the mission

(Fig. 4). Most images are captured in the afternoon (after driving), with some panoramas captured in the late morning (mid-drive), and all Martian seasons are represented (Fig. 5). There are typically five stereo pairs per panorama, covering a full 360 around the rover, but some panoramas only capture three or four pairs. The image resolution is typically 2560x1920, but some images are 1280x960 and 5120x3840. For consistency, we down-sample all images to 1280x960 resolution.

For ground truth locations, an independent team of mapping experts produced rover orthomosaics and hand-aligned them to the HiRISE orbital map. The accuracy of these locations is expected to be within a couple of pixels of the orbital map (0.5m) [28]. This is a mature process that the Mars missions currently depend on for localization, and our results comport with this expected accuracy.

### B. Benchmark Results on Mars Panoramas

We use a 30m search range, a conservative bound on the rover odometry position uncertainty after a long drive. We label localizations with  $>5m$  error as outliers based on a Perseverance mission requirement for global localization.

Our approach achieves 0.36m accuracy across 264 panoramas with no significant outliers (Table I) and a maximum error of 1.07m. Notably, the mode is around one-pixel error (0.25m) in the orbital map (Fig. 7). Other common approaches, including masked normalized cross-correlation (NCC) and mutual information, produce some accurate localizations but suffer from significant outlier rates, especially with large search ranges. The modified census transform exhibits greater invariance to the intensity differences between the rover images and the orbital images and universally produces less ambiguous correlation peaks in the score matrix compared to alternative approaches (Fig. 8).

### C. Ablation & Sensitivity Studies

Table II is an ablation study showing the relative impact of each pipeline stage on overall localization performance. The modified census transform is the key contributor to localization performance. Radiometric correction, subpixel estimation, DEM matching, and stereo method individually play a more minor role in localization performance, but in combination they impact average error and outlier rate enough to justify their inclusion.

Table III is a sensitivity study showing the relative impact of the key parameters. A key performance factor is maximizing the orthomosaic coverage, either through increasing the number of panorama stereo pairs around the rover, or increasing the max stereo range. The approach is surprisingly insensitive to lower image resolutions, likely due to the comparatively low resolution of the orbital map (Fig. 3). It is also insensitive to increasing search ranges or map resolution, although run-time and accuracy will be impacted in obvious ways for each.

### D. Results on the Perseverance Rover

The proposed framework has been tested both on the Perseverance Mars Rover on Mars, and a flight-identical

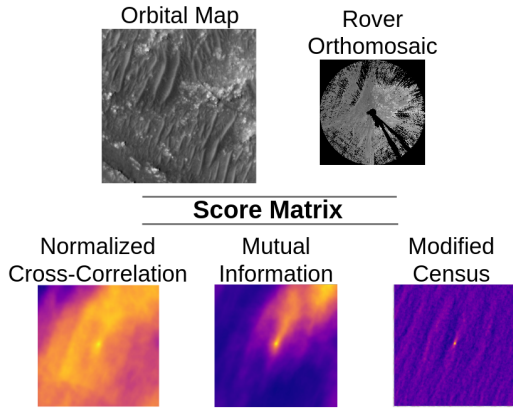


Fig. 8: Compared to normalized cross-correlation (NCC) and mutual information, the modified census transform consistently produces the least ambiguous correlation peaks.

TABLE I: Comparison to Alternative Similarity Measures

Algorithm	Avg. Error	99%-Tile	>5m Errors
Masked NCC	24.97m	45.90m	225/264
Masked Mutual Information	24.33m	46.98m	216/264
Modified Census (Ours)	<b>0.36m</b>	<b>0.93m</b>	<b>0/264</b>

TABLE II: Ablation Study

Removed Feature	Avg. Error	99%-Tile	>5m Errors
Baseline (No removals)	0.36m	0.93m	0/264
No Radiometric Correction	0.38m	0.94m	1/264
No Subpixel Estimation	0.37m	0.93m	0/264
No SGM Stereo (Local BM)	0.36m	0.91m	0/264
No DEM Scores	0.36m	0.93m	0/264
No RadCorr/Subpx/SGM/DEM	0.73m	11.70m	3/264
No Census Scores	14.17m	43.26m	185/264

version of the rover on Earth. On both rovers, the algorithm runs on a Qualcomm Snapdragon 801 processor, an onboard co-processor to the rover’s main RAD750 computer.

On the flight-identical rover on Earth, the algorithm was benchmarked on the Snapdragon co-processor: for 5 stereo pairs, it ran in 81s and used <1GB of volatile memory (Table IV). Semi-global stereo matching represents the majority of the run-time and memory consumption, although a more constrained platform could use a simpler local stereo method without a significant performance impact (Table III).

On Mars, the framework was deployed on the Perseverance rover and successfully executed on Martian sol 914. The system executed in shadow mode, using images from a previous sol without updating rover pose. The small footprint of our system was key to enabling this test, due to the limited uplink bandwidth to Mars: the binary was 744KB, with a small 250KB appearance and elevation orbital map representing a 200m x 200m area.

## V. CONCLUSION AND FUTURE WORK

Motivated by increasing autonomous rover navigation, this paper discusses our developments towards enabling longer

TABLE III: Sensitivity Study

Category	Parameter	Avg. Error	99%-Tile	>5m Errors
Number Stereo Pairs	3 pairs max	0.38m	1.04m	0/264
	2 pairs max	0.42m	1.11m	0/264
	1 pair max	0.88m	20.09m	6/264
Max Stereo Range	40m	0.36m	0.93m	0/264
	20m	0.35m	0.92m	0/264
	10m	0.69m	11.11m	3/264
Stereo Resolution	1280x960	0.36m	0.93m	0/264
	640x480	0.35m	0.91m	0/264
	320x240	0.35m	0.91m	0/264
	160x120	0.40m	1.01m	1/264
	80x60	4.17m	35.31m	46/264
Search Radius	30m	0.36m	0.93m	0/264
	50m	0.36m	0.93m	0/264
	100m	0.36m	0.93m	0/264
Appearance Map Resolution	150m	0.92m	0.95m	1/264
	25cm	0.36m	0.93m	0/264
	28cm	0.36m	0.94m	0/264
Appearance Map Resolution	50cm	0.48m	1.02m	0/264
	100cm	0.71m	1.72m	0/264

TABLE IV: Real execution timing on flight-identical hardware. At conservative thermal settings, the algorithm runs within two minutes and <1GB volatile memory.

Resolution	Pairs	Total	Stereo	Matching	Memory
2560x1920	1	36s	14s	8s	<1GB
2560x1920	5	81s	65s	8s	<1GB

drives applied to the Perseverance rover on Mars. The approach meets human level accuracy. It enables the rover to be commanded to drive for potentially unlimited drive distances without requiring localization from Earth. Censible is a system for global localization, applicable to wide range of planetary robots.

Beyond Perseverance, absolute position estimation is key for future planetary robotic missions. A number of national agencies and private companies have plans for developing planetary rovers in the next decade. Lunar rover missions like Endurance aim to traverse many kilometers across the Moon’s south pole. The Ingenuity and future Sample Retrieval Helicopters require absolute position estimation to navigate autonomously and retrieve sample tubes. Censible’s success has already begun to impact some of these mission studies.

## VI. ACKNOWLEDGEMENTS

This work is supported by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004). We thank the Ingenuity helicopter and Perseverance rover teams, including Gerik Kubiak, Andrei Tumar, Jeffrey Biesadecki, Yang Cheng, Bob Deen, Oleg Parisier, Francois Ayoub, Joseph Bowkett, and Mark Maimone for their contributions to this work.

## REFERENCES

- [1] V. Verma, M. W. Maimone, D. M. Gaines, R. Francis, T. A. Estlin, S. R. Kuhn, G. R. Rabideau, S. A. Chien, M. M. McHenry, E. J. Graser, A. L. Rankin, and E. R. Thiel, "Autonomous robotics is driving perseverance rover's progress on mars," *Science Robotics*, vol. 8, no. 80, 2023. [Online]. Available: <https://www.science.org/doi/abs/10.1126/scirobotics.adi3099>
- [2] T. Parker, M. Golombek, and M. Powell, "Geomorphic/geologic mapping, localization, and traverse planning at the opportunity landing site, mars," in *41st Annual Lunar and Planetary Science Conference*, no. 1533, 2010, p. 2638.
- [3] T.-H. Pham, W. Seto, S. Daftry, B. Ridge, J. Hansen, T. Thrush, M. Van der Merwe, G. Maggolino, A. Brinkman, J. Mayo, *et al.*, "Rover relocation for mars sample return by virtual template synthesis and matching," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 4009–4016, 2021.
- [4] B. Van Pham, A. Maligo, and S. Lacroix, "Absolute map-based localization for a planetary rover," in *12th Symposium on Advanced Space Technologies and Automation in Robotics*, 2013, pp. 1–8.
- [5] D. Geromichalos, M. Azkarate, E. Tsardoulias, L. Gerdes, L. Petrou, and C. Perez Del Pulgar, "Slam for autonomous planetary rovers with global localization," *Journal of Field Robotics*, vol. 37, no. 5, pp. 830–847, 2020.
- [6] A. Ansar and L. Matthies, "Multi-modal image registration for localization in titan's atmosphere." IEEE, 2009, pp. 3349–3354.
- [7] D. Gaines, G. Doran, M. Paton, B. Rothrock, J. Russino, R. Mackey, R. Anderson, R. Francis, C. Joswig, H. Justice, *et al.*, "Self-reliant rovers for increased mission productivity," *Journal of Field Robotics*, vol. 37, no. 7, pp. 1171–1196, 2020.
- [8] R. Brockers, P. Proença, J. Delaune, J. Todd, L. Matthies, T. Tzanetos, and J. B. Balaram, "On-board absolute localization based on orbital imagery for a future mars science helicopter," in *2022 IEEE Aerospace Conference (AERO)*. IEEE, 2022, pp. 1–11.
- [9] L. Matthies, S. Daftry, S. Tepsuporn, Y. Cheng, D. Atha, R. M. Swan, S. Ravichandar, and M. Ono, "Lunar rover localization using craters as landmarks," *arXiv preprint arXiv:2203.10073*, 2022.
- [10] R. Li, K. Di, A. B. Howard, L. Matthies, J. Wang, and S. Agarwal, "Rock modeling and matching for autonomous long-range mars rover localization," *Journal of Field Robotics*, vol. 24, no. 3, pp. 187–203, 2007.
- [11] E. Boukas, A. Gasteratos, and G. Visentin, "Introducing a globally consistent orbital-based localization system," *Journal of Field Robotics*, vol. 35, no. 2, pp. 275–298, 2018.
- [12] R. Li, K. Di, J. Wang, S. Agarwal, L. Matthies, A. Howard, and R. Willson, "Incremental bundle adjustment techniques using networked overhead and ground imagery for long-range autonomous mars rover localization," in *Proc. of The 8th International Symposium on Artificial Intelligence, Robotics and Automation in Space—ISAIRAS, ESA SP-603*, 2005.
- [13] K. Di, F. Xu, J. Wang, S. Agarwal, E. Brodyagina, R. Li, and L. Matthies, "Photogrammetric processing of rover imagery of the 2003 mars exploration rover mission," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 63, no. 2, pp. 181–201, 2008.
- [14] S. Chiodini, M. Pertile, S. Debei, L. Bramante, E. Ferrentino, A. G. Villa, I. Musso, and M. Barrera, "Mars rovers localization by matching local horizon to surface digital elevation models," in *2017 IEEE International Workshop on Metrology for AeroSpace (MetroAeroSpace)*. IEEE, 2017, pp. 374–379.
- [15] K. Ebadi, K. Coble, D. Kogan, D. Atha, R. Schwartz, C. Padgett, and J. Vander Hook, "Toward autonomous localization of planetary robotic explorers by relying on semantic mapping," in *2022 IEEE Aerospace Conference (AERO)*. IEEE, 2022, pp. 1–10.
- [16] —, "Semantic mapping in unstructured environments: Toward autonomous localization of planetary robotic explorers," in *Proceedings of the 2022 IEEE Aerospace Conference, Big Sky, MT, USA, 2022*, pp. 5–12.
- [17] T. Ahmad, G. Bebis, M. Nicolescu, A. Nefian, and T. Fong, "Horizon line detection using supervised learning and edge cues," *Computer Vision and Image Understanding*, vol. 191, p. 102879, 2020.
- [18] F. Cozman, E. Krotkov, and C. Guestrin, "Outdoor visual position estimation for planetary rovers," *Autonomous Robots*, vol. 9, no. 2, pp. 135–150, 2000.
- [19] F. Cozman and E. Krotkov, "Automatic mountain detection and pose estimation for teleoperation of lunar rovers," in *Experimental Robotics V*. Springer, 1998, pp. 207–215.
- [20] J. Vander Hook, R. Schwartz, K. Ebadi, K. Coble, and C. Padgett, "Topographical landmarks for ground-level terrain relative navigation on mars," in *2022 IEEE Aerospace Conference (AERO)*. IEEE, 2022, pp. 1–6.
- [21] S. Wang, Y. Zhang, A. Vora, A. Perincherry, and H. Li, "Satellite image based cross-view localization for autonomous vehicle," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 3592–3599.
- [22] V. Franchi and E. Ntigiou, "Planetary rover localisation via surface and orbital image matching," in *2022 IEEE Aerospace Conference (AERO)*. IEEE, 2022, pp. 1–14.
- [23] I. M. i Caireta, "Improving global localization algorithms for mars rovers with neural networks," 2021.
- [24] A. Bal and H. Palus, "Image vignetting correction using a deformable radial polynomial model," *Sensors*, vol. 23, no. 3, p. 1157, 2023.
- [25] D. B. Gennery, "Generalized camera calibration including fish-eye lenses," *International Journal of Computer Vision*, vol. 68, no. 3, pp. 239–266, 2006.
- [26] H. Hirschmuller, "Stereo processing by semiglobal matching and mutual information," *IEEE Transactions on pattern analysis and machine intelligence*, vol. 30, no. 2, pp. 328–341, 2007.
- [27] R. Zabih and J. Woodfill, "Non-parametric local transforms for computing visual correspondence," in *Computer Vision—ECCV'94: Third European Conference on Computer Vision Stockholm, Sweden, May 2–6 1994 Proceedings, Volume II 3*. Springer, 1994, pp. 151–158.
- [28] R. Deen. (2022) Mars 2020 project rover places data products for pds. Accessed on 2023-10-3. [Online]. Available: [https://pds-geosciences.wustl.edu/m2020/urn-nasa-pds-mars2020\\_rover\\_places/document/Mars2020\\_Rover\\_PLACES\\_PDS\\_SIS.pdf](https://pds-geosciences.wustl.edu/m2020/urn-nasa-pds-mars2020_rover_places/document/Mars2020_Rover_PLACES_PDS_SIS.pdf)