

# MARTIAN: A Rendering Framework for Aerial Mars Imagery from HiRISE Orbital Data

Dario Pisanti<sup>1,\*</sup>, Georgios Georgakis<sup>2</sup>

**Abstract**—Aerial navigation on Mars requires vision-based pipelines that are robust to the diverse illumination conditions and terrain morphology of the Martian surface. A key bottleneck for training and evaluating such methods is the scarcity of large-scale, annotated aerial datasets. We present MARTIAN, an open-source Blender-based rendering framework that leverages real HiRISE orbital map products to synthesize realistic aerial views of the Martian terrain under controllable lighting conditions and at varying altitudes. MARTIAN generates observations with accurate pose annotations, directly addressing the scarcity of training data for vision-based navigation on Mars. The framework has been validated through its deployment in concurrent work on map-based localization systems for Ingenuity and future Mars rotorcraft, where synthetically trained deep image matchers were successfully evaluated on real Mars imagery. MARTIAN is publicly available at: <https://github.com/nasa-jpl/martian>.

## I. INTRODUCTION

The demonstration flights of NASA’s Mars Helicopter, *Ingenuity*, have marked a groundbreaking milestone in Mars exploration [1]. One of the next evolutionary steps is the Mars Science Helicopter (MSH) [2], a hexacopter concept with multi-kg payload capacity and 10 km traverse capability at altitudes up to 100 m, aimed at supporting high-priority investigations in Martian astrobiology, climate, and geology [3]. Similarly, the Skyfall mission aims to deploy six scout helicopters mid-air in the Martian atmosphere from the entry capsule, thus eliminating the need for a landing platform [4]. The deployed UAVs will independently perform exploration tasks for resource prospecting, science measurements, and safe landing site inspection for future crewed missions.

These aerial platforms will need to perform precision onboard vision-based navigation and localization to enable long-range flights over uncertain terrain in a global navigation satellite system (GNSS)-denied environment. Images captured by the onboard navigation camera are typically processed by a Visual Inertial Odometry (VIO) system to produce position estimates in a relative navigation fashion. However, VIO drift needs to be bounded over long traverses with an absolute localization strategy. This can be accomplished by registering onboard imagery onto orbital maps pre-registered to a global reference frame and estimating the navigation camera pose relative to the map. This inherently drift-free geo-localization technique is referred to as Map-based Localization (MbL).

Both VIO and MbL methods rely on the identification of distinctive features in onboard imagery and orbital maps to enable relative and absolute navigation, respectively. This process poses significant perception challenges due to diverse illumination conditions and terrain morphology, and is even more pronounced in MbL due to large differences in lighting and scale between onboard imagery and the reference map. This problem is also relevant for vision-based Terrain Relative Navigation during the Entry, Descent and Landing (EDL) phase of Martian landers[5], which rely on orbital maps taken at restricted times of day.

While recent deep learning methods [6], [7] have demonstrated robustness to illumination and scale variations on in-the-wild datasets such as MegaDepth [8] for terrestrial applications, the main bottleneck is the lack of large-scale datasets that would allow finetuning these methods in planetary domains. Unlike Earth-based applications relying on abundant data [9], large-scale and annotated aerial datasets are not readily available for Mars. Imagery data from the Mars2020 mission provide aerial observations from the Landing Vision System (LVS) camera during EDL and from the navigation camera onboard *Ingenuity*. However, the lack of accurate pose annotations prevents these data to be used for a comprehensive robustness study on scale and illumination variations. Existing large-scale Martian datasets, such as A14MARS [10], address surface-based navigation by providing terrain classification labels for rover driving, while the SynMars-Air dataset [11] addresses aerial object detection and terrain segmentation from a helicopter viewpoint.

To better support learning-based perception methods for aerial navigation on Mars, we developed the Mars Aerial Rendering Tool for Imaging and Navigation (MARTIAN). This simulation framework leverages real map products from the Mars Reconnaissance Orbiter (MRO) High-Resolution Imaging Science Experiment (HiRISE) [12] to generate large-scale datasets for training vision-based navigation pipelines. MARTIAN has been deployed in concurrent work, as detailed in Section III.

## II. MARTIAN FRAMEWORK

We take advantage of real map products created from HiRISE to create a large-scale dataset suitable for training vision-based pipelines for aerial navigation on Mars. We developed a python-based framework in the open-source 3D computer graphics software Blender to import HiRISE Digital Terrain Models (DTMs) and textured ortho-projected images to generate maps and aerial observations of a Martian site under various lighting conditions and at

<sup>1</sup>Space Robotics Research Group, SnT, University of Luxembourg.

<sup>2</sup>Jet Propulsion Laboratory, California Institute of Technology.

\*Corresponding author: [dario.pisanti@uni.lu](mailto:dario.pisanti@uni.lu).

different altitudes.

**HiRISE data.** With its resolution capability at nadir of 25 centimeters per pixel from 300 km altitude, HiRISE has been serving as an indispensable orbital asset for identifying and selecting safe landing sites for robotics exploration missions. In this work, we utilize a 1 m / post DTM and a 0.25 m / pixel ortho-image with equirectangular projection generated from stereo pairs imaging of the Jezero Crater landing site for the Mars2020 Mission, over an area of 6.737 km by 14.403 km.

**Terrain and texture modeling.** The Jezero HiRISE DTM was imported in Blender 4.0 using a modified version of the original HiRISE import plug-in [13]. This add-on can load the terrain data at a desired resolution within a range (0., 100] % of the full model resolution and generate a mesh. By leveraging terrain metadata, including geographic information and resolution, a UV map is created for ortho-images to be draped over the companion DTM and used as a high-fidelity terrain texture. The MARTIAN framework initially imports the Jezero DTM with 10% of its original resolution, allowing for efficient terrain setup. Then, a material is added to the mesh, with its surface shading model being controlled via a Principled Bidirectional Scattering Distribution Function (BSDF). The 0.25 m / pixel Jezero ortho-image is then loaded in the shading editor as the base texture for the terrain surface. The ortho-image coordinates data are retrieved by a Texture Coordinate Node to ensure that the texture is properly mapped onto the 3D mesh. Finally, the mesh is reloaded with its full resolution of 1 m / post.

**Scene setup and camera modeling.** MARTIAN allows for setting multiple scenes for perspective and orthographic imaging with user-defined camera intrinsics and extrinsics. Figure 1 shows a view of the Jezero crater site in MARTIAN. Given the *world* reference frame  $W$  defined as an East-North-Up coordinate system with origin on the terrain map center, the camera can be located in  $W$  by providing the  $xy$ -coordinates and the altitude (in meters) with respect to the terrain at those coordinates. To position the camera object above the terrain mesh at the desired altitude, MARTIAN adopts a ray-tracing approach by using a Bounding Volume Hierarchy tree, a data structure used by Blender to efficiently organize geometric objects in 3D space. The camera frame  $C$  is centered at the optical center, with its  $X$ -axis pointing to right along the image width, the  $Z$ -axis pointing towards the terrain, and the  $Y$ -axis completing the orthogonal set. The attitude of the  $C$  with respect to the world frame  $W$  is specified as input, and the final pose ( $\mathbf{R}_{WC}|\mathbf{t}_{WC}$ ) is saved, where  $\mathbf{t}_{WC}$  is the location of the camera in world coordinated and  $\mathbf{R}_{WC}$  is the rotation matrix that aligns  $W$  to  $C$ .

**Lighting.** MARTIAN provides the capability to render scenes in Blender with various illumination conditions

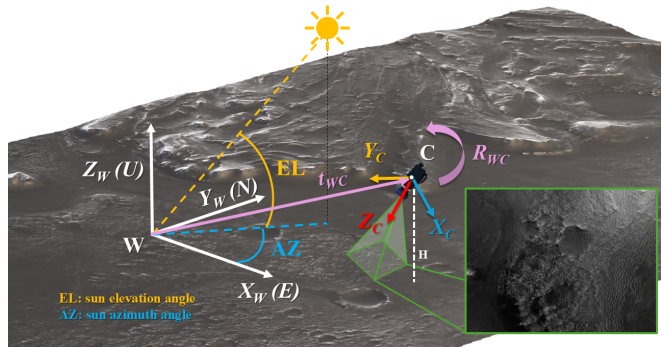


Fig. 1. View of the Jezero crater site in MARTIAN with the simulated perspective imaging camera and a rendered observation.

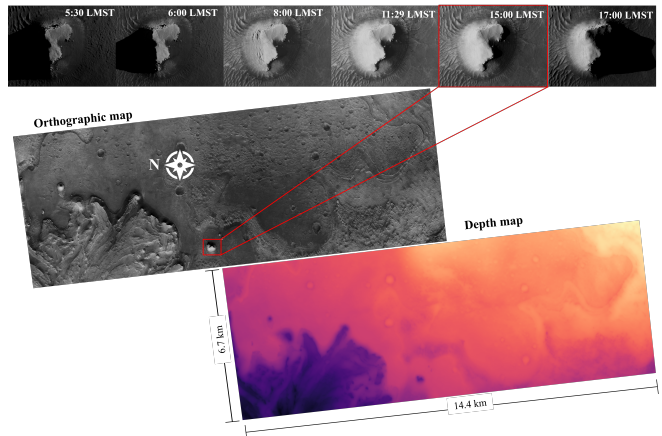


Fig. 2. Orthographic grayscale map rendered in MARTIAN at 15:00 Local Mean Solar Time (LMST) on the Jezero crater site on 2031-05-10, with the corresponding depth map. Simulated aerial observations from the same Martian day are shown from 05:30 to 17:00 LMST.

by tuning Sun light source parameters such as irradiance, apparent disk diameter and orientation in the map frame. Higher irradiance values cast brighter illumination and shadows in the scene, while the angular diameter of the Sun disk as seen by the scene controls the softness versus harshness of the shadows. Sun orientation simulating a specific time of day is specified by the user through elevation (EL) and azimuth angles (AZ) as shown in Figure 1. Lighting computations are performed using the Blender Cycles engine. This is a physically-based rendering engine that uses a ray tracing algorithm to accurately simulate light behavior. Sun irradiance is set to the maximum value of 590 W/m<sup>2</sup> and angular diameter to 0.35°, consistent with actual estimations for Mars. The BSDF parameters and camera exposure were tuned empirically to match the visual appearance of the real ortho-projected map. Figure 2 illustrates the effect of different times of day on Mars on a rendered aerial observation from the Jezero crater site.

### III. APPLICATIONS IN AERIAL MAP-BASED LOCALIZATION

MARTIAN was used to generate a large-scale synthetic aerial dataset to train deep image matchers for the

autonomous Map-Based Localization (MbL) of Mars rotorcraft. The dataset comprises 17 orthographic grayscale maps rendered from combinations of Sun azimuth ( $0^\circ - 360^\circ$ ,  $45^\circ$  steps) and elevation ( $30^\circ$ ,  $60^\circ$ ,  $90^\circ$ ), one corresponding orthographic depth map, and 4500 nadir-pointing aerial observations at fixed Sun angles  $AZ=180^\circ$  and  $EL=40^\circ$  with their corresponding depth images. Observations were randomly sampled from the HiRISE Jezero Crater DTM with uniform distribution in the  $(x, y)$  world-frame coordinates and within the altitude range  $[64, 200]$  m., where the 64 m lower bound reflects MARTIAN’s resolution limit, tied to the HiRISE texture resolution. Camera extrinsics and altitude data were stored for each observation and used as ground truth.

**MbL for Ingenuity [14].** This dataset was used to pre-train LoFTR [6], a detector-free image matching method that uses Transformers [15] in a coarse-to-fine strategy to produce pixel-wise semi-dense correspondences. The pre-trained model was then incorporated into an MbL system for Ingenuity, with a training strategy designed to address the scarcity of annotated in-domain data. Since only a limited number of Ingenuity navigation images are available and obtaining accurate pose annotations is particularly challenging due to noisy VIO estimates and frequent low-texture terrain, the approach first leverages MARTIAN data to obtain a suitable initial representation for HiRISE map features, before fine-tuning on a small set of Ingenuity images with self-supervised pseudo annotations. This intermediate pre-training on rendered data proved beneficial: fine-tuning LoFTR directly on Ingenuity images without it resulted in an 11.2%  $Acc@5m$  drop across all flights, and a gap of up to 26.4% on the most challenging terrain. The MbL pipeline achieved 89.4%  $Acc@5m$  and near-perfect 99.8%  $Acc@10m$  across Ingenuity flights, outperforming prior template-matching and hand-crafted feature methods, while generalizing to previously unseen flights with as few as 30 training images from a single flight.

**MbL for future Mars rotorcraft under challenging lighting [16].** The same dataset was used to train Geo-LoFTR, a geometry-aided extension of LoFTR that incorporates geometric context from DTMs derived from orbital data products into the feature matching process. Geo-LoFTR was adopted in an MbL pipeline targeting the long-range navigation of future Mars rotorcraft at altitudes up to 200 m, such as the Mars Science Helicopter (MSH)[2], demonstrating improved localization accuracy over prior methods by up to 31.8%  $Acc@1m$  under challenging illumination conditions. In particular, the geometric context was shown to increase robustness to varying Sun elevation and azimuth angles, as well as to the scale difference between onboard observations and the reference map, where Geo-LoFTR maintained consistent performance across the full altitude range. Robustness was also demonstrated over a full simulated Martian day at the Jezero crater site, from 6:00 to 17:00 Local Mean Solar Time (LMST), against a HiRISE-like reference map captured

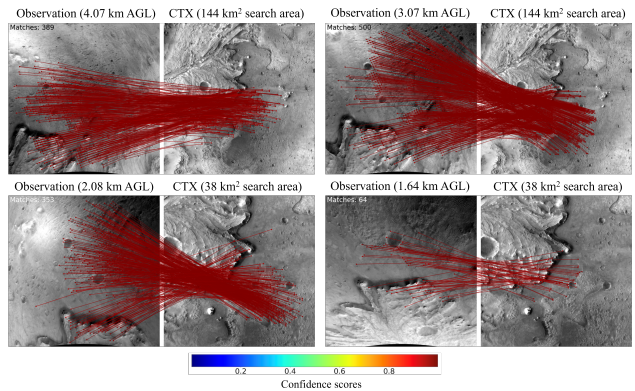


Fig. 3. Geo-LoFTR matches between CTX map crops and LCAM observations at different altitudes Above Ground Level (AGL) from the imaging sequence during the Mars2020 EDL.

at 15:00 LMST. Furthermore, Geo-LoFTR was validated on real Mars2020 descent imagery, where frames captured by the onboard Lander Vision System camera (LCAM)[5] were matched against a Context Camera (CTX) map [17] at 6 m/px resolution. Despite being trained entirely on MARTIAN-rendered data for the anticipated flight altitude of MSH, Geo-LoFTR successfully provided global localization solutions across descent frames in the 6 km -960 m altitude range, indicating preliminary synthetic-to-real transfer under comparable observation-to-map scale ratios. Figure 3 shows an example of Geo-LoFTR matches between LCAM frame captured at different altitudes and CTX map crops obtained by leveraging pose priors. However, broader generalization across terrain types and map sources still needs further investigation.

#### IV. CONCLUSION

We presented MARTIAN, a Blender-based rendering framework that leverages HiRISE map products to generate large-scale annotated aerial datasets for Mars, addressing a critical gap for training vision-based navigation pipelines. MARTIAN has already demonstrated utility in concurrent work on map-based localization for both Ingenuity and future Mars rotorcraft. A current limitation is that HiRISE textures are ortho-images captured at a specific time of day, and therefore already embed the illumination and shadow conditions of that acquisition, which the renderer cannot fully decouple. Additionally, while preliminary validation of MARTIAN-trained methods on real Mars imagery suggests a promising sim-to-real transfer, further studies across diverse landing sites are needed to fully characterize generalization.

#### REFERENCES

- [1] H. F. Grip, D. Conway, J. Lam, N. Williams, M. P. Golombek, R. Brockers, M. Mischna, and M. R. Cacan, “Flying a helicopter on mars: How ingenuity’s flights were planned, executed, and analyzed,” in *2022 IEEE Aerospace Conference (AERO)*, 2022, pp. 1–17.
- [2] T. Tzanetos, J. Bapst, G. Kubiak, L. P. Tosi, S. Sirlin, R. Brockers, J. Delaune, H. F. Grip, L. Matthies, J. Balam, S. Withrow-Maser, W. Johnson, and B. Pipenberg, “Future of mars rotorcraft - mars science helicopter,” in *IEEE Aerospace Conference*, 2022, pp. 1–16.

- [3] J. Bapst, T. J. Parker, J. Balaram, T. Tzanetos, L. H. Matthies, C. D. Edwards, A. Freeman, S. Withrow-Maser, W. Johnson, E. Amador-French, J. L. Bishop, I. J. Daubar, C. M. Dundas, A. A. Fraeman, C. W. Hamilton, C. Hardgrove, B. Horgan, C. W. Leung, Y. Lin, A. Mittelholz, and B. P. Weiss, "Mars Science Helicopter: Compelling Science Enabled by an Aerial Platform," *Bulletin of the AAS*, vol. 53, no. 4, mar 18 2021, <https://baas.aas.org/pub/2021n4i361>.
- [4] NASA, "SkyFall Mission – Ingenuity Mars Helicopter," <https://science.nasa.gov/mission/mars-2020-perseverance/ingenuity-mars-helicopter/>, Mar. 2026, accessed: April 2026.
- [5] A. E. Johnson, Y. Cheng, N. Trawny, J. F. Montgomery, S. Schroeder, J. Chang, D. Clouse, S. Aaron, and S. Mohan, "Implementation of a map relative localization system for planetary landing," *Journal of Guidance, Control, and Dynamics*, vol. 46, no. 4, pp. 618–637, 2023. [Online]. Available: <https://doi.org/10.2514/1.G006780>
- [6] J. Sun, Z. Shen, Y. Wang, H. Bao, and X. Zhou, "Loftir: Detector-free local feature matching with transformers," in *IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 8922–8931.
- [7] J. Edstedt, Q. Sun, G. Bökman, M. Wadenbäck, and M. Felsberg, "Roma: Robust dense feature matching," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 19 790–19 800.
- [8] Z. Li and N. Snavely, "Megadepth: Learning single-view depth prediction from internet photos," in *IEEE conference on computer vision and pattern recognition*, 2018, pp. 2041–2050.
- [9] R. Zhu, L. Yin, M. Yang, F. Wu, Y. Yang, and W. Hu, "Sues-200: A multi-height multi-scene cross-view image benchmark across drone and satellite," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 33, no. 9, pp. 4825–4839, 2023.
- [10] R. M. Swan, D. Atha, H. A. Leopold, M. Gildner, S. Oij, C. Chiu, and M. Ono, "Ai4mars: A dataset for terrain-aware autonomous driving on mars," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 1982–1991.
- [11] Y. Qi, X. Xiao, M. Yao, Y. Xiong, L. Zhang, and H. Cui, "Airformer: Learning-based object detection for mars helicopter," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 18, pp. 100–111, 2025.
- [12] A. S. McEwen, E. M. Eliason, J. W. Bergstrom, N. T. Bridges, C. J. Hansen, W. A. Delamere, J. A. Grant, V. C. Gulick, K. E. Herkenhoff, L. Keszthelyi, R. L. Kirk, M. T. Mellon, S. W. Squyres, N. Thomas, and C. M. Weitz, "Mars reconnaissance orbiter's high resolution imaging science experiment (hirise)," *Journal of Geophysical Research: Planets*, vol. 112, no. E5, 2007.
- [13] PhaseIV, "Blender hirise dtm importer," 2024, version 1.0. [Online]. Available: <https://github.com/phaseIV/Blender-Hirise-DTM-Importer>
- [14] G. Georgakis, D. Pisanti, N. Williams, C. Mauceri, G. Kubiak, A. Ansar, and R. Brockers, "A map-based localization system for ingenuity using deep image matching," *IEEE Transactions on Field Robotics*, vol. 2, pp. 787–804, 2025.
- [15] A. Waswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *NIPS*, 2017.
- [16] D. Pisanti, R. Hewitt, R. Brockers, and G. Georgakis, "Vision-based geo-localization of future mars rotorcraft in challenging illumination conditions," *arXiv preprint arXiv:2502.09795*, 2025.
- [17] M. C. Malin, J. F. Bell III, B. A. Cantor, M. A. Caplinger, W. M. Calvin, R. T. Clancy, K. S. Edgett, L. Edwards, R. M. Haberle, P. B. James *et al.*, "Context camera investigation on board the mars reconnaissance orbiter," *Journal of Geophysical Research: Planets*, vol. 112, no. E5, 2007.