

High-fidelity 3D reconstruction for planetary exploration

A. Martínez-Petersen¹, L. Gerdes¹, D. Rodríguez-Martínez¹, C. J. Pérez-del-Pulgar¹

Abstract—Planetary exploration increasingly relies on autonomous robotic systems capable of perceiving, interpreting, and reconstructing their surroundings in the absence of global positioning or real-time communication with Earth. Rovers operating on planetary surfaces must navigate under severe environmental constraints, limited visual redundancy, and communication delays, making onboard spatial awareness and visual localization key components for mission success. Traditional techniques based on Structure-from-Motion (SfM) or Simultaneous Localization and Mapping (SLAM) provide geometric consistency but struggle to capture radiometric detail or to scale efficiently in unstructured, low-texture terrains typical of extraterrestrial environments. This work explores the integration of radiance field-based methods—specifically Neural Radiance Fields (NeRF) and Gaussian Splatting—into a unified, automated environment reconstruction pipeline for planetary robotics. Our system combines the *Nerfstudio* and *COLMAP* frameworks with a Robot Operating System 2 (ROS 2)-compatible workflow capable of processing raw rover data directly from *rosvbag* recordings. This approach enables the generation of dense, photorealistic, and metrically consistent 3D representations from minimal visual input, supporting improved perception and planning for autonomous systems operating in planetary-like conditions. The resulting pipeline establishes a foundation for future research in radiance field-based mapping, bridging the gap between geometric and neural representations in planetary exploration.

I. INTRODUCTION

Planetary exploration is rapidly evolving towards missions of increasing complexity and duration. In this context, rovers have become the main agents of observation and operation, adapting to increasingly extreme conditions over extended periods of time and reducing the costs and risks associated with direct human intervention in missions.

However, the challenges faced by rovers are numerous. Due to communication delays, absence of GNSS, and limited environmental knowledge, rover operators must rely on a combination of orbital maps, local images from the rover’s perspective, and the rover’s own pose estimation, which inherently contains errors. Given the possibility of losing a mission due to the failure to detect an obstacle, operations must be planned conservatively, making it essential to improve the autonomy and situational awareness of both the planning tools and the robot itself.

To enhance such autonomy in the absence of reliable absolute references, current localization and perception systems in planetary exploration rely on multisensor data sources such as stereo cameras, Inertial Measurement Units (IMUs),

wheel encoders, laser rangefinders, and, in the near future, solid-state LiDAR sensors, to estimate relative motion and reconstruct the immediate environment. Techniques such as visual odometry allow the estimation of the trajectory and the generation of local maps by detecting and tracking features between successive stereo image pairs. However, the accuracy of these estimates can be severely limited under specific conditions: low-texture or homogeneous terrains (e.g., flat sand or uniform rock surfaces), large illumination changes (e.g., deep shadows or strong sunlight), low overlap between consecutive frames due to fast rover motion, and repetitive or ambiguous geometric patterns. Under these circumstances, visual correspondence becomes unreliable, causing pose errors to rapidly accumulate. Some proposed alternatives, such as Simultaneous Localization and Mapping (SLAM), face significant limitations in space missions due to their high energy consumption, sensitivity to harsh environmental conditions, and the low visual redundancy of rover paths, often resulting in sparse or noisy models that fail to capture the radiometric information of the environment adequately.

By contrast, recent 3D reconstruction approaches jointly integrate geometry, texture, and volumetric density into a single model, enabling the generation of high-fidelity, photorealistic reconstructions even from limited visual data. Unlike traditional geometric methods, continuous radiance-based representations are able to synthesize novel views, interpolate between camera positions, and capture complex lighting effects, opening new possibilities for visual perception and autonomous planning in planetary environments. Their direct use in robotic platforms remains constrained, however, by practical barriers, such as the need for extensive preprocessing, complex parameter tuning, and incompatibility with rover-specific data formats and sensor structures, which hinder their deployment beyond controlled research settings.

The main contributions of this work are summarized as follows:

- The design of a fully automated and modular pipeline for planetary environment reconstruction integrating *Nerfstudio* [1] and *COLMAP* [2], [3] within a unified workflow.
- Direct processing of raw ROS2 *rosvbag* [4] recordings, including multisensor synchronization, metadata enrichment with orientation priors, and automated format conversion for radiance field reconstruction.
- An enhanced geometric initialization strategy based on an externalized and updated *COLMAP* Structure-from-Motion (SfM) stage, improving robustness and convergence under exploration-like conditions.

¹A. Martínez-Petersen, L. Gerdes, C. J. Pérez-del-Pulgar, and D. Rodríguez-Martínez are with the Institute for Research in Mechatronics Engineering and Cyber-Physical Systems (IMECH) of the University of Málaga, Andalucía Tech, 29070 Málaga, Spain (e-mail: david.rm@uma.es).

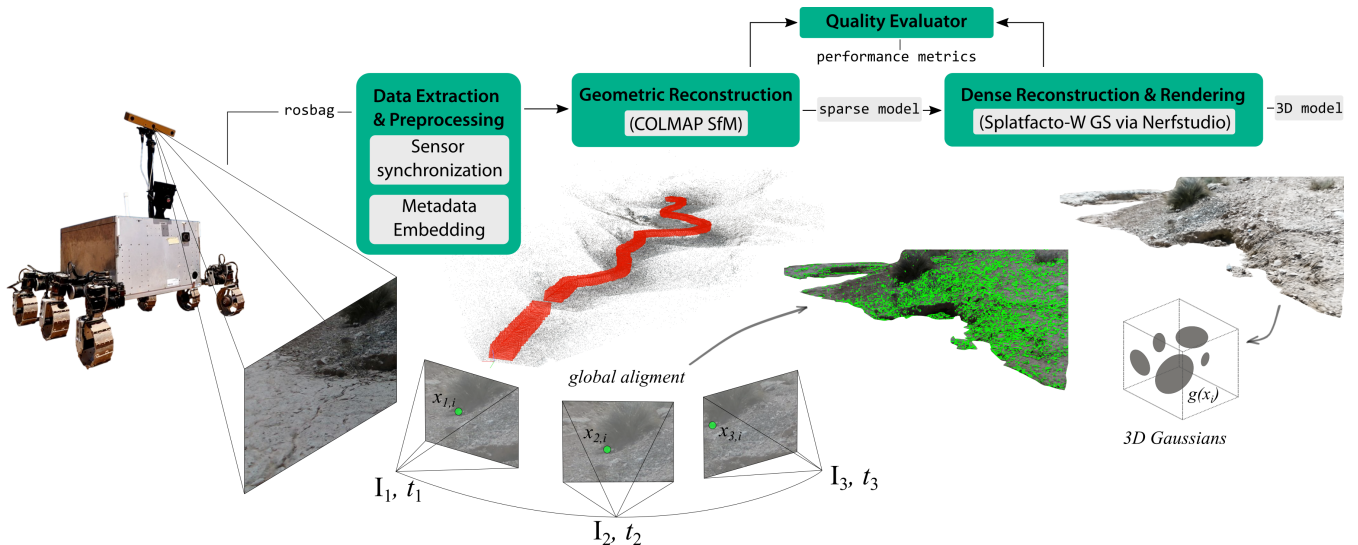


Fig. 1. Overview of the proposed reconstruction pipeline. Raw rover data (such as RGB images, IMU readings, and GNSS data) are extracted directly from *rosbag* recordings and preprocessed for geometric reconstruction with *COLMAP* (Structure-from-Motion (SfM)). The resulting sparse model and camera poses are then passed to the radiance field stage, where *Splatfacto-W* performs dense reconstruction through Gaussian Splatting, generating a photorealistic 3D model. Finally, a quality evaluation module computes performance metrics across both geometric and visual dimensions.

- A containerized implementation ensuring reproducibility, portability, and practical deployment across different computing platforms and operating systems.

II. RADIANCE FIELD-BASED METHODS

A. Neural Radiance Fields (NeRF)

NeRF are a 3D reconstruction technique that employs deep neural networks to continuously represent a three-dimensional scene from a set of 2D images [5]. The model learns a function that maps a 3D position \mathbf{x} and a viewing direction $\mathbf{d} = (\theta, \phi)$ (defined in spherical coordinates, where θ is the polar angle and ϕ the azimuthal angle) to a volumetric density $\sigma(\mathbf{x})$ and color $\mathbf{c}(\mathbf{x}, \mathbf{d}) = (R, G, B)$, as follows:

$$\mathbf{C}(\mathbf{r}) = \int_{t_1}^{t_2} T \sigma \mathbf{c} dt \quad (1)$$

where $\mathbf{C}(\mathbf{r})$ is the expected color observed along camera ray $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$, and $T(t)$ is the accumulated transmittance along the ray.

In its discrete form:

$$\mathbf{C}(\mathbf{r}) = \sum_{i=1}^N \alpha_i T_i \mathbf{c}_i, \quad T_i = \exp\left(-\sum_{j=1}^{i-1} \sigma_j \delta_j\right), \quad (2)$$

where i indexes the sampled points along each ray, δ_i is the distance between consecutive samples, α_i denotes their opacity and $\mathbf{c}_i = (R_i, G_i, B_i)$ the corresponding color.

Essentially, NeRF is a model that learns how light and color exist in a scene. It uses a neural network to figure out what color each pixel should be by tracing rays of light from each camera pixel and learning the brightness and density of every point in space.

Several NeRF variants address similar challenges to those encountered in the 3D reconstruction of planetary environments. The original NeRF [5] achieves very high-fidelity reconstructions but is computationally demanding. NeRF in the Wild (NeRF-W) [6] adds robustness to variable lighting and uncalibrated images, common in planetary exploration. Nerfacto [1] improves training efficiency, enabling the reconstruction of large scenes on resource-limited platforms. These variants exemplify the trade-offs between accuracy, robustness, and efficiency for rover trajectories.

B. Gaussian Splatting

In contrast, the Gaussian Splatting method does not rely on neural networks but on an explicit representation based on 3D points associated with Gaussian distributions. Each 3D scene point $\mathbf{x} \in \mathbb{R}^3$ is modeled as a Gaussian defined by its mean $\boldsymbol{\mu}$, covariance matrix Σ , and amplitude α :

$$G(x) = \alpha \exp\left[-\frac{1}{2}(x - \boldsymbol{\mu})^T \Sigma^{-1} (x - \boldsymbol{\mu})\right]. \quad (3)$$

During rendering, these 3D Gaussians are projected onto the 2D image plane according to the camera parameters. The projection is computed by transforming the covariance matrix from world coordinates to image space as:

$$\Sigma' = JW\Sigma W^T J^T, \quad \Sigma' \in \mathbb{R}^{2 \times 2} \quad (4)$$

where W represents the world-to-camera transformation matrix, and J is the Jacobian of the perspective projection that maps 3D points onto the 2D image plane. This operation determines the apparent shape and spread of each Gaussian when viewed through the camera, producing the corresponding 2D *splats* used in differentiable rendering.

Unlike NeRF, which implicitly learns the scene through neural networks, Gaussian Splatting directly employs a set

of 3D primitives (or *Gaussians*), optimizing their position, color, size, and opacity parameters from multiple images. The result is a visually coherent reconstruction with lower computational cost and faster convergence times. While Neural Radiance Fields (NeRF) produces highly detailed models at the expense of intensive training, 3D Gaussian Splatting achieves a more efficient and explicitly parameterized representation that is easier to update, making it particularly suitable for time-constrained or resource-limited applications [7]. Variants of the original method as Splatfacto [7], improve Gaussian Splatting by jointly optimizing the Gaussians for both efficiency and reconstruction quality. Its extension, Splatfacto in the Wild (Splatfacto-W) [8], further improves robustness to varying lighting conditions and uncalibrated camera data, similar in spirit to NeRF-W.

C. Comparison of Methods Applied to Space Robotics

Selecting the most suitable radiance field approach for planetary rover data requires balancing visual fidelity, computational efficiency, and robustness to challenging acquisition conditions. The comparison presented in Table I does not aim to be exhaustive but rather to justify the selection of the most appropriate model for our proposed pipeline. Six representative variants were selected for their relevance to planetary exploration scenarios, which are characterized by irregular trajectories, limited image overlap, variable illumination, and constrained onboard computing resources. The parameters reported in Table I capture complementary aspects of performance: while processing time, GPU memory consumption, and reconstruction accuracy describe the expected trade-offs between efficiency and quality, robustness to lighting and calibration reflects how well each method tolerates real-world imaging imperfections. Finally, the suitability for rover trajectories highlights each method’s capacity to handle the motion patterns typical of planetary rovers, such as long straight segments, low parallax, and sudden viewpoint changes, which challenge conventional radiance field reconstruction pipelines.

Within NeRF-based methods, the original NeRF provides the reference for maximum visual quality, NeRF-W introduces robustness to uncalibrated and uncontrolled data, and Nerfacto optimizes the trade-off between performance and reconstruction accuracy. On the other hand, Gaussian-based methods such as 3D Gaussian Splatting, Splatfacto, and Splatfacto-W offer explicit, faster, and more flexible representations better suited for real-world and resource-limited applications.

Given the characteristics of rover datasets—raw RGB images, variable-velocity rover motion, and strong illumination variability—*Splatfacto-W* was identified as the most appropriate method. This variant combines the computational efficiency of Gaussian Splatting with the robustness mechanisms of NeRF-W, enabling stable reconstruction under uncontrolled conditions while maintaining low memory requirements and high adaptability to non-uniform trajectories.

Consequently, it was selected as the reconstruction core of the proposed pipeline.

III. METHODOLOGY

A. Pipeline Description

The proposed pipeline automates the complete 3D reconstruction workflow, from raw rover data to fully rendered visual models through a series of modular terminal scripts that ensure full reproducibility, scalability, and minimal user intervention. Each stage is independently executable yet seamlessly integrated into a unified framework, standardizing parameter configurations and guaranteeing consistent performance across experiments and hardware setups.

The pipeline, as described in Figure 1, is organized into four main stages, each corresponding to a key phase in the results analysis:

- 1) **Data Extraction and Preprocessing:** Raw mission data are extracted directly from Robot Operating System 2 (ROS 2) *rosbags*. This stage automates the selection of RGB image sequences, the import of intrinsic and extrinsic camera parameters, and the registration of GNSS or odometry-based coordinates when available. The resulting dataset is formatted for geometric reconstruction and can optionally include localization priors to guide subsequent stages.
- 2) **Geometric Reconstruction (*Structure-from-Motion*):** Unlike the default *Nerfstudio* integration, this pipeline employs an external instance of *COLMAP*, configured with extended functionality for pose initialization and prior-based optimization. Different *preprocessing configurations* can be applied—such as *sequential matching*, which restricts feature matching to temporally adjacent frames, *exhaustive matching*, which tests all image pairs, or the proposed hybrid method combining sequential matching with camera priors—to optimize correspondence search and convergence stability. The output includes camera trajectories, reconstructed 3D points, and geometric quality metrics such as reprojection error, number of feature observations, and track length.
- 3) **Dense Reconstruction and Rendering:** The output of *COLMAP* is automatically converted into the format required by *Nerfstudio*, where dense 3D reconstruction is performed based on *Splatfacto-W*. This step produces high-fidelity photorealistic models that can be interactively visualized and rendered from novel viewpoints, allowing both qualitative inspection and quantitative evaluation of visual consistency.
- 4) **Quantitative and Qualitative Evaluation:** The final stage extracts and aggregates performance metrics to assess both geometric and visual quality. Geometric consistency is quantified through reprojection error, number of feature observations per image, and mean track length, while rendering fidelity is evaluated using perceptual image metrics such as Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index Measure

TABLE I
COMPARISON OF NERF AND GAUSSIAN SPLATTING VARIANTS RELEVANT TO ROVER-BASED 3D RECONSTRUCTION.

Parameter	NeRF	NeRF-W	Nerfacto	3DGS	Splatfacto	Splatfacto-W
Processing time	Slow	Slow	Medium	Fast	Very fast	Fast
Reconstruction accuracy	Very high	High	High	High	High	Medium–High
Memory consumption	Very high	Very high	High	High	Medium	Medium
Robustness to lighting / calibration	Low	High	Medium	High	High	Very High
Suitability for rover trajectories	Low	Medium	Medium	Medium	High	Very High

(SSIM), and Learned Perceptual Image Patch Similarity (LPIPS). These indicators provide a comprehensive view of reconstruction accuracy, radiometric realism, and overall system robustness.

All stages are orchestrated within a unified execution environment to ensure reproducibility and efficiency. To enhance portability, the complete pipeline is containerized using Docker, integrating dedicated environments for *COLMAP*, *Nerfstudio*, and ROS 2, together with all dependencies. This design allows seamless deployment across different operating systems and computational infrastructures, including local workstations and GPU clusters. The modular structure also facilitates the integration of alternative reconstruction backends or sensor modalities, extending the framework’s applicability to different mission scenarios.

B. Experimental Setup

The pipeline was executed on a high-performance server, equipped with multiple NVIDIA A6000 GPUs, high-speed NVMe storage, and a 10 Gbps NAS system, enabling real-time visualization and the reconstruction of large-scale datasets. This infrastructure ensures efficient data management, high throughput, and the capability to process extended rover trajectories that exceed the memory limits of standard desktop systems.

C. Pipeline Validation

The proposed workflow was validated using real rover data from the Bardenas Semi-Desert Planetary Rover Dataset (BASEPROD) dataset [9], acquired by the Planetary Robotics Laboratory (PRL) of the European Space Agency (ESA) and the Space Robotics Lab (SRL) of the University of Malaga (UMA) using the Martian Rover Testbed for Autonomy (MaRTA) rover [10]. The experiments were carried out in the Bardenas Reales semi-desert (Navarra, Spain), a Mars-analog site characterized by heterogeneous textures, irregular terrain geometry, and highly variable natural illumination. Although the data were not originally captured for photogrammetric purposes—resulting in motion blur and non-optimized viewpoints—these conditions make the dataset a realistic testbed for assessing reconstruction performance under rover-like constraints. The full dataset, including both the extracted data and the original ROS 2 *rosbags*, is publicly available in ESA’s robotics datasets repository [11].

Five representative segments were selected from different traverses to ensure diversity in visual appearance, trajectory type, and image count. Trajectories T1 through T5 in this paper correspond to segments of traverses 2023-07-22_14-18-23, 2023-07-21_12-38-15, 2023-07-23_13-05-11, 2023-07-23_12-52-39, and 2023-07-21_17-34-18 in BASEPROD. As we needed a different format than the already extracted data for this application, we extracted the required data directly from the *rosbags* using the automated data extraction stage of the pipeline.

Each segment was processed independently through the complete workflow, comprising data extraction, geometric reconstruction with *COLMAP*, and dense visual modeling using *Splatfacto-W*.

The experiments were conducted under outdoor lighting with variable solar incidence angles and included slight deviations in rover speed and camera orientation. This configuration provides a realistic approximation of planetary rover operations, testing the robustness of the proposed method against real-world noise, imperfect calibration, and heterogeneous imaging conditions.

IV. RESULTS

This section summarizes the outcomes of applying the proposed pipeline to the rover datasets from the BASEPROD dataset, as described in Section III. Both qualitative and quantitative results are presented to assess reconstruction accuracy, visual fidelity, and computational efficiency.

A. Overview of the Reconstructions

Representative examples of the reconstructed scenes are shown in Figure 2 and Figure 3. These visual results illustrate the capability of the proposed pipeline to recover coherent 3D structure and radiometric appearance from real-world rover data acquired under uncontrolled outdoor conditions. The models correspond to Mars-analog terrains in the Bardenas Reales site, characterized by heterogeneous textures, irregular topography, and variable natural illumination.

The aerial views in Figure 2 highlight the overall surface continuity and topographic consistency achieved in the reconstructed models, while Figure 3 compares real camera images with their corresponding renderings. The visual alignment between the two domains demonstrates the geometric and photometric consistency attained through the Gaussian Splatting approach implemented in *Splatfacto-W*.



Fig. 2. General aerial views of 3D models reconstructed from rover traverses in BASEPROD. Each model corresponds to a segment of the terrain captured by the MaRTA rover under real outdoor illumination and surface conditions.



Fig. 3. Comparative examples between real-world views (left) and their corresponding 3D reconstructions (right). The visual alignment demonstrates the high geometric and radiometric consistency achieved by the proposed workflow.

B. Performance Metrics

To quantitatively evaluate the proposed pipeline, both computational performance and reconstruction quality were analyzed using the metrics outlined in Table II and Table III. These indicators quantify the efficiency of preprocessing and the geometric and photometric accuracy of the resulting pipeline.

Table II compares the execution times of three preprocessing configurations: (1) the default *Nerfstudio* method, (2) a sequential image-matching approach optimized for ordered rover trajectories, and (3) the proposed configuration, which combines sequential matching with GNSS-based camera priors. In GNSS-denied scenarios, these priors could be replaced by camera pose estimations computed through

Visual Odometry (VO) or any other available estimation. The **Match [s]** column measures the time required for feature correspondence between images, while **Map [min]** indicates the duration of the mapping stage, where camera poses and sparse 3D structures are estimated. The **Conv.** column denotes whether the reconstruction converged successfully. Together, these values reflect the computational cost and robustness of the geometric preprocessing pipeline.

Table III focuses on the quality of the resulting reconstructions. The **Reprojection Error** expresses the average pixel deviation between the reprojected 3D points and their actual image locations (a lower value indicates higher geometric accuracy). The **Avg. Obs.** (average number of observations per image) represents how many distinct 3D points are typi-

TABLE II

COMPARISON OF RECONSTRUCTION TIMES FOR DIFFERENT PREPROCESSING CONFIGURATIONS: NERFSTUDIO DEFAULT, SEQUENTIAL MATCHING, AND OUR PROPOSED METHOD WITH SEQUENTIAL MATCHING AND CAMERA PRIORS.

Trajectory	Nerfstudio (Default)			Sequential Matching			Proposed Method (Seq. + Priors)		
	Match [min]	Map [min]	Conv.	Match [min]	Map [min]	Conv.	Match [min]	Map [min]	Conv.
T1	6.18	68.1	Yes	0.447	56.37	Yes	0.441	26.858	Yes
T2	5.32	50.3	No	0.353	25.00	Yes	0.354	19.354	Yes
T3	13.8	143.85	Yes	1.2	54.564	Yes	1.2	51.00	Yes
T4	6.53	96.57	No	0.366	51.197	No	0.36	22.995	Yes
T5	6.25	73.6	No	0.297	40.88	Yes	0.319	25.725	Yes

cally visible in each frame, which relates to the redundancy and overlap of the input data. **Track Len.** measures the mean number of images in which each 3D point is observed, assessing temporal coherence across the trajectory.

To complement the geometric assessment, visual fidelity was measured using three standard rendering metrics: **PSNR** (Peak Signal-to-Noise Ratio) quantifies pixel-level similarity between rendered and real images; **SSIM** (Structural Similarity Index Measure) evaluates perceived structural consistency; and **LPIPS** (Learned Perceptual Image Patch Similarity) approximates perceptual quality from a human-vision standpoint. For these metrics, higher PSNR and SSIM values and lower LPIPS scores indicate more realistic visual reconstructions.

C. Results Evaluation and Limitations

Overall, the developed workflow successfully reconstructed all trajectories with consistent geometric and visual quality. Three algorithmic configurations were compared: (1) the default *Nerfstudio* preprocessing, (2) a sequential matching implementation, and (3) our proposed method combining sequential matching with GNSS-based priors. The latter achieved the best results, reducing average matching and mapping times by 93.1% and 65.8%, respectively, while maintaining full reconstruction convergence. This demonstrates the benefits of integrating sequential matching and prior motion information when dealing with ordered, trajectory-based image datasets.

From the geometric perspective, the SfM evaluation yielded low mean reprojection errors (below 0.6 px), high average observations per image (above 5,000), and track lengths exceeding 7 frames in most cases, confirming the internal consistency and robustness of COLMAP’s reconstruction. These results indicate sufficient overlap and visual redundancy, even in scenes with limited texture.

Quantitative evaluation of the radiance field rendering was performed using PSNR, SSIM, and LPIPS metrics. Across the five trajectories, average PSNR values ranged between 26 and 29 dB, SSIM between 0.78 and 0.83, and LPIPS between 0.20 and 0.26, denoting acceptable-to-good perceptual quality. Among all the trajectories evaluated, T5 achieved the best results. This appears to be due to its highly regular and almost rectilinear capture pattern, which improves feature tracking and results in a more stable and well-conditioned

bundle adjustment. However, a limited viewpoint diversity reduces the geometric coverage of the scene, worsening the quality of the reconstruction when observed from previously unseen angles. Despite the use of training views for evaluation, which potentially inflates absolute scores, the results consistently reflect coherent color reproduction, geometric alignment, and visually plausible renderings.

A qualitative inspection of the reconstructions corroborates these findings: reconstructed models exhibit accurate geometry and continuous surfaces, though thin objects and vegetation remain challenging. Some regions near the image boundaries exhibit ghosting and aliasing artifacts, typical of low-overlap sequences. Nevertheless, the temporal coherence between consecutive rendered frames is satisfactory, which is essential for navigation and visualization applications.

In summary, the proposed pipeline achieved reliable 3D reconstructions under Mars-analog conditions, balancing visual fidelity, computational efficiency, and robustness. Experiments were conducted on five representative segments from the BASEPROD dataset, all acquired at a single Mars-analog site. Although these segments exhibit variability in terrain texture, illumination conditions, and rover trajectory geometry, cross-site validation on additional planetary-analog datasets is still required to fully assess generalization. Nevertheless, the current results confirm the viability of the approach for rover-based scene reconstruction and perception. However, several limitations of the current workflow must be acknowledged. Despite its robustness, the pipeline still relies on offline processing and requires computational resources far beyond those available on present-day planetary rover platforms, preventing real-time or onboard deployment. Performance also depends heavily on COLMAP’s SfM stage, which constitutes the main bottleneck of the workflow: although the results reported above demonstrate good average consistency, sparse reconstruction remains sensitive to texture richness, viewpoint overlap, and lighting variability. Any degradation in COLMAP’s initialization or feature matching propagates downstream and directly affects reconstruction quality. Importantly, Gaussian Splatting is particularly vulnerable to limited viewpoint diversity. Predominantly frontal or straight-line trajectories often provide insufficient parallax for stable splat optimization, reducing the model’s ability to generalize to novel viewpoints even when sparse geometry is

TABLE III
 INTEGRATED EVALUATION OF RECONSTRUCTION QUALITY ACROSS ALL TRAJECTORIES. COLOR CODING FOLLOWS THE LEGEND BELOW,
 EMPHASIZING HIGHER-QUALITY RESULTS IN GREEN TONES.

Trajectory	Reproj. Err. [px]	Avg. Obs. [#]	Track Len. [#]	PSNR [dB]	SSIM	LPIPS
T1	0.552	6437.78	7.07	28.2 ± 3.0	0.806 ± 0.116	0.231 ± 0.119
T2	0.548	6287.21	6.98	27.1 ± 1.8	0.803 ± 0.048	0.204 ± 0.041
T3	0.541	5594.43	7.43	28.2 ± 3.0	0.805 ± 0.116	0.231 ± 0.119
T4	0.520	6141.99	6.91	26.1 ± 3.5	0.783 ± 0.107	0.264 ± 0.121
T5	0.476	5310.43	7.92	29.3 ± 3.9	0.827 ± 0.133	0.224 ± 0.101
Excellent	<0.2 px	> 6000	> 7	>30 dB	> 0.9	< 0.15
Good	0.2–0.5 px	4000–6000	5–7	28–30 dB	0.8–0.9	0.15–0.25
Acceptable	0.5–1.0 px	2000–4000	3–5	25–28 dB	0.7–0.8	0.25–0.35
Poor	>1.0 px	< 2000	< 3	<25 dB	< 0.7	> 0.35

nominally acceptable. Another practical limitation may arise from inherent ROS2 image acquisition pipeline irregularities as inconsistent frame rates or timestamping reduce the effectiveness of sequential matching, negatively affecting both SfM and radiance field optimization.

V. CONCLUSION

This work presented a fully automated and modular pipeline for 3D reconstruction from rover data, integrating a version of *COLMAP* with prior-based structure estimation and dense reconstruction using Gaussian Splatting within *Nerfstudio*. The pipeline is designed to handle ROS 2-based datasets, standardize preprocessing, and produce high-quality photorealistic models, while providing quantitative metrics for both geometric and visual evaluation. We validated our proposed workflow on real rover data from the BASEPROD dataset, demonstrating reliable reconstructions across trajectories. Quantitative results show that incorporating sequential matching and pose priors reduces matching and mapping times by 93.1% and 65.8%, respectively, compared to *Nerfstudio*'s default configuration, while improving reconstruction accuracy and convergence.

While our evaluation focused on preprocessing configurations and internal pipeline optimization, a direct end-to-end comparison against alternative mapping approaches (e.g., SLAM) and radiance-field implementations remains an important direction for future studies.

Although the modified process achieved successful reconstructions in Mars-analog environments, it still faces limitations in speed, robustness, and data dependency. The pipeline currently operates offline, preventing real-time mapping or onboard decision-making, and relies on *COLMAP* for structure estimation, which makes performance sensitive to initialization, texture richness, viewpoint overlap, and lighting variability. These factors can lead to inconsistent results, especially in low-texture or repetitive terrains, and should, therefore, guide future improvements.

Overall, this work confirms the feasibility and potential of these technologies, providing a documented and operational foundation for future research in radiance-field-based spatial reconstruction using rover platforms.

ACKNOWLEDGMENTS

This work was part of the project INSIGHT (PID2024-160373OB-C21) funded by MICIU / AEI / 10.13039/501100011033 / FEDER, UE. It was partially supported by the European Space Agency under activity no. 4000140043/22/NL/GLC/ces. It was also technically supported by IMECH.UMA through PPRO-IUI-2023-02.

REFERENCES

- [1] M. Tancik, E. Weber, E. Ng, R. Li, B. Yi, J. Kerr, T. Wang, A. Kristoffersen, J. Austin, K. Salahi, A. Ahuja, D. McAllister, and A. Kanazawa, "Nerfstudio: A modular framework for neural radiance field development," in *ACM SIGGRAPH 2023 Conference Proceedings*, ser. SIGGRAPH '23, 2023.
- [2] J. L. Schönberger and J.-M. Frahm, "Structure-from-motion revisited," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [3] J. L. Schönberger, E. Zheng, M. Pollefeys, and J.-M. Frahm, "Pixel-wise view selection for unstructured multi-view stereo," in *European Conference on Computer Vision (ECCV)*, 2016.
- [4] S. Macenski, T. Foote, B. Gerkey, C. Lalancette, and W. Woodall, "Robot operating system 2: Design, architecture, and uses in the wild," *Science Robotics*, vol. 7, no. 66, p. eabm6074, 2022.
- [5] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "NeRF: Representing scenes as neural radiance fields for view synthesis," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020.
- [6] R. Martin-Brualla, N. Radwan, M. S. M. Sajjadi, J. T. Barron, A. Dosovitskiy, and D. Duckworth, "NeRF in the Wild: Neural radiance fields for unconstrained photo collections," 2021.
- [7] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis, "3D Gaussian splatting for real-time radiance field rendering," *ACM Transactions on Graphics*, vol. 42, no. 4, pp. 1–13, 2023.
- [8] C. Xu, J. Kerr, and A. Kanazawa, "Splatfacto-W: A nerfstudio implementation of Gaussian splatting for unconstrained photo collections," *arXiv preprint arXiv:2407.12345*, 2024.
- [9] L. Gerdes, T. Wiese, R. C. Arquillo, L. Bielenberg, M. Azkarate, H. Leblond, F. Wilting, J. O. Cortés, A. Bernal, S. Palanco, and C. P. del Pulgar, "BASEPROD: The Bardenas semi-desert planetary rover dataset," *Scientific Data*, vol. 11, no. 1, p. 1054, Sep. 2024.
- [10] M. Azkarate, L. Gerdes, T. Wiese, M. Zwick, M. Pagnamenta, J. Hidalgo-Carrio, P. Poulakis, and C. J. P. del Pulgar, "Design, testing, and evolution of Mars rover testbeds: European Space Agency planetary exploration," *IEEE Robotics & Automation Magazine*, vol. 29, pp. 10–23, 9 2022.
- [11] L. Gerdes, R. Arquillo-Castilla, T. Wiese, L. Bielenberg, M. Azkarate, H. Leblond, F. Wilting, J. Ortega-Cortés, A. Bernal, S. Palanco, and C. Pérez del Pulgar, "BASEPROD Data Repository," 8 2024. [Online]. Available: <https://doi.org/10.57780/esa-xxd1ysw>