Lunar Surface Reconstruction via Neural Rendering of KPLO LUTI Imagery

Dunam Kim[1], Suwan Lee[1], Kibaek Park[1], Jo Ryeong Yim[2], Dong-Gyu Kim[2], Eunhyeuk Kim[2], Seokju Lee[1] (slee@kentech.ac.kr), [1]Korea Institute of Energy Technology (KENTECH), [2]Korea Aerospace Research Institute (KARI), South Korea.

Introduction: High-precision digital elevation models (DEMs) of the Moon are essential for mission planning and scientific analysis. However, traditional photogrammetric methods often struggle to maintain spatial resolution and consistency when processing pushbroom camera imagery, which generates vast, continuous image lines. Furthermore, because most lunar images are acquired in a nadir (straight-down) orientation to support cartographic and geological research, limited parallax—that is, minimal or ambiguous differences in viewing angles between images—makes it difficult to derive accurate 3D structure from standard stereo matching. To address these challenges, we propose a custom neural rendering framework based on a simplified Neural Radiance Field (NeRF) [1]. While NeRF typically requires multiple consistent images from different viewpoints for high-fidelity 3D rendering, our approach simulates multi-view geometry by treating each pushbroom image “line” as a distinct viewpoint and modeling the spacecraft’s orbit as a continuous path along the planetary surface—all from a single pass. We also show that our framework remains robust even when fewer pushbroom lines are available, making it particularly promising for large-scale reconstructions.

Method: Our approach trains a deep neural network to represent the lunar surface as an implicit volumetric rendering function, using a pushbroom camera pose and the corresponding ray directions as input. Each line of the pushbroom camera is treated as a distinct viewpoint: we sample 3D points along these rays, pass them through the network, and obtain densities and grayscale intensities. These outputs are then integrated to form radiance and metric-scale depth, guided by SLDEM to ensure accurate surface elevations.

A key aspect of our framework is multi-view training via sequential pushbroom lines. As the spacecraft travels along the lunar surface, any desired viewpoint can be approximated by continuing along its flight path—eliminating the need to distinguish between single and multi-shot scenarios. Since many pushbroom camera poses are noisy or interpolated, we optionally incorporate learnable pose estimation and surface priors (e.g., smoothness losses, local radius constraints) to improve geometric consistency.

Unlike conventional DEM approaches [2, 3, 4] that rely on explicit disparity maps or stereo matching, our method iteratively refines the scene geometry by backpropagating pixel intensity and depth errors, capturing subtle parallax and sensor geometry effects often missed by classical photogrammetry. Through a differentiable volume rendering equation, the network implicitly learns a continuous 3D representation of the lunar terrain—both spatial density and view-dependent radiance. During forward rendering, the model calculates contributions from sampled depths along each ray to produce the final grayscale radiance. The discrepancies between these rendered intensities and observed data are then used to update the volumetric surface representation, steadily improving reconstruction accuracy.

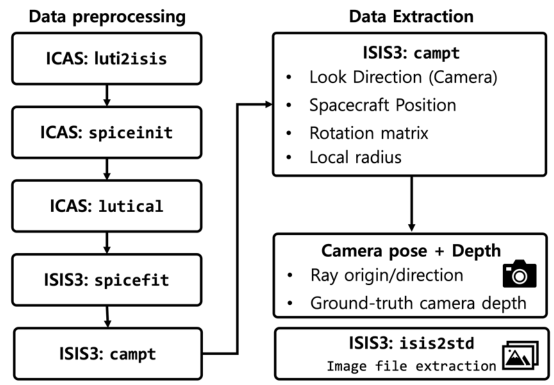

Figure 1. Data Preprocessing Pipeline. ICAS-based applications attach kernel data to the LUTI images and perform radiometric calibration. Next, ISIS3 tools extract photometry and geometry data for training our 3D reconstruction AI model. The resulting ray information, transformation matrices, and ground-truth depths are then used as inputs for the neural networks.

Dataset: The Lunar Terrain Imager (LUTI) on board the Korea Pathfinder Lunar Orbiter (KPLO) features a focal length of 404 ± 0 mm and achieves a pixel scale of 2.5 m/pixel at an altitude of 100 km. We evaluate our approach on Tycho Crater (43.37°S, 348.68°E) captured by LUTI which offers relatively high-quality pushbroom imagery and sufficient LOLA 3D point coverage, enabling reliable quantitative and qualitative analyses. In our dataset, we use SLDEM as the reference depth for metric-scale training because it provides dense global coverage (±60° in latitude, 360° in longitude) without empty pixels and is accurately co-registered to LOLA points, ensuring consistent geometry and minimal alignment errors. All images in our dataset are captured at small phase angles as possible for minimizing the differences in surface

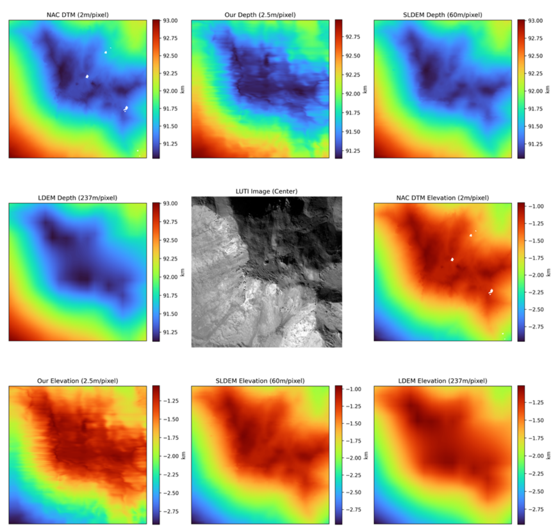

Figure 2. Comparison of depth and elevation results from our model and DEM. A 3×3 grid visualization of depth and elevation maps. The center shows the ground-truth image, with depth maps placed on the left and top, and elevation maps on the right and bottom. Our method preserves the depth distribution of existing DEMs while capturing more detailed lunar surface features in both depth and elevation representations.

observations. As shown in Figure 1, we apply a data preprocessing pipeline that provides the essential inputs for our 3D reconstruction model.

Experiments: We compare our approach with existing DEMs by evaluating vertical accuracy, pixel resolution, and the number of images required for DEM generation. In Table 1, LOLA 3D points from the test region are projected onto both the baseline DEMs and our reconstruction to compute mean squared vertical errors. Notably, our method achieves high-quality reconstructions from a single image, whereas baseline DEMs require multiple stereo pairs or additional LOLA points (see the 4-th column of Table 1.) This single-shot learning capability readily generalizes to various existing DEMs, because our framework applies weak supervision—guiding the network to capture coarse depth distributions rather than overfitting to precise ground-truth labels.

In Figure 2, while preserving SLDEM's overall depth distribution, our approach refines coarse depths to match the resolution of the input imagery, yielding significantly sharper surface details. Crucially, it avoids overfitting to either shadowed pixels or over-smoothed depths by learning a robust surface representation from combined geometric and photometric cues. This robustness extends to subtle features often lost due to ambiguous parallax—minimal or unclear viewpoint differences that make it difficult to derive accurate 3D structure solely from pixel shifts.

Method	Pixel Scale (m/pixel)	Vertical Accuracy (m)	Configuration of Tycho crater
NAC DTM [2]	2	2.82	8 NAC stereo pairs (32 images)
SLDEM [3]	59	4.66	SELENE DEM + LOLA points
LDEM [4]	236.9	23.58	LOLA points with interpolation
Ours	2.5	66.1	1 LUTI image

Table 1. Quantitative comparison of depth and elevation results with the baseline DEMs (NAC DTM, SLDEM, and LDEM). Note that all baseline DEMs are interpolated to a 2.5 m/pixel.

Discussion & Limitation: We propose a novel neural volumetric rendering method for generating higher-quality lunar DEMs, robust to camera-pose noise, shadows, and the Moon's unique reflectance. This framework can be applied to any planetary imagery captured by pushbroom cameras, alongside their associated DEMs. Most notably, our method continuously produces a high-resolution DEM at the original image scale without requiring additional interpolation. However, our method still relies on metric depth from SLDEM, even though it internally learns an enhanced 3D representation. Future work will investigate more sophisticated approaches that remove the need for metric depth labels. Additionally, the final pixel scale and geometric sharpness are tied to the resolution of the input imagery, suggesting a multi-scale strategy as a potential avenue for further improvement.

Acknowledgments: We would like to extend our sincere gratitude to the Korea Aerospace Research Institute (KARI) and the Korea Astronomy and Space Science Institute (KASI) for their invaluable advice and comments throughout this research endeavor. This research was supported by the Korea Astronomy and Space Science Institute under the R&D program (Project No. 2024-1-904-03) supervised by the Ministry of Science and ICT.

References:

[1] B. Mildenhall, et al., "NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis." European Conference on Computer Vision (ECCV), 2020.

[2] Henriksen, M. R., et al. "Extracting accurate and precise topography from LROC narrow angle camera stereo observations." *Icarus* 283 (2017): 122-137.

[3] Barker, M. K., et al. "A new lunar digital elevation model from the Lunar Orbiter Laser Altimeter and SELENE Terrain Camera." *Icarus* 273 (2016): 346-355.

[4] Smith, David E., et al. "The lunar orbiter laser altimeter investigation on the lunar reconnaissance orbiter mission." *Space science reviews* 150 (2010): 209-241.