

LNEM: Lunar Neural Elevation Model

Suwan Lee¹ Jo Ryeong Yim² Kibaek Park¹ Dong-Gyu Kim²
Eunhyeuk Kim² Minsup Jeong³ Chae Kyung Sim³ Seokju Lee^{1†}

¹KENTECH ²KARI ³KASI

<https://viewlab-group.github.io/LNEM/>




Figure 1. **Lunar Studio Benchmark and representative LNEM reconstructions.** Lunar Studio provides a standardized multi-orbit dataset integrating NAC and LUTI observations, while LNEM leverages rigorous pushbroom sensor modeling to produce geometrically consistent and high-fidelity elevation reconstructions. The results show clear preservation of terrain contours and stable reconstruction quality across diverse lunar regions and illumination conditions.

Abstract

High-resolution and high-precision digital elevation models (DEMs) of the lunar surface are essential for landing site selection and geological research. However, traditional stereo matching provides a limited representation of the 3D scene and struggling with non-textured regions and extreme illumination variations. Recent lunar neural rendering methods are also ill-suited for 3D reconstruction due to their reliance on simple pinhole approximations for pushbroom sensors. These challenges are further compounded by geometric misalignment, distributional bias, and labor-intensive handcrafted preprocessing in satellite image pipelines. To address these issues, we introduce the Lunar Neural Elevation Model (LNEM), a volumetric reconstruction method that explicitly incorporates the pushbroom imaging process. A core component of our approach is Lunar Studio, a multi-orbit dataset and pipeline constructed using Rigorous Sensor Models (RSMs) to produce geometrically consistent observations from the Lu-

nar Reconnaissance Orbiter Camera (LROC) Narrow Angle Camera (NAC) and the Korea Pathfinder Lunar Orbiter (KPLO) Lunar Terrain Imager (LUTI). LNEM integrates this pushbroom camera formulation with learned shadow modeling, enabling geometrically grounded and illumination-aware volumetric rendering under challenging lunar lighting conditions. Extensive experiments demonstrate that LNEM achieves geometrically consistent reconstruction across multiple sensors under diverse viewing and illumination conditions, providing a scalable complement to conventional DEM pipelines. To support reproducibility and future lunar research, we release Lunar Studio, the multi-orbit dataset, and the LNEM reconstruction pipeline.

1. Introduction

One of the most critical tasks in planetary exploration is generating accurate digital elevation models (DEMs) from satellite imagery. In the context of lunar missions, DEMs are pivotal for landing site selection, rover navigation, and geological mapping, where high-fidelity reconstructions of-

[†]Corresponding author.

fer deeper insights into the origin and geological evolution of the lunar surface.

Conventional photogrammetric pipelines rely on stereo matching between overlapping image tiles [6], but often struggle with non-textured regions and the pushbroom geometry of satellite sensors, where images are captured line-by-line and the camera pose changes with each scan. This geometry precludes straightforward epipolar constraints and complicates correspondence search [6, 11]. Furthermore, illumination variation across orbits violates brightness consistency assumptions of traditional stereo, and terrain misalignments commonly require laser altimeter corrections (*e.g.*, via LOLA) for reliable absolute elevations.

While extensive high-resolution lunar imagery exists from years of orbital capture, acquiring a geometrically consistent dataset suitable for high-fidelity neural rendering remains difficult due to complex, non-linear orbital trajectories and drastic altitude variations. To circumvent this, current methods often resort to using only a sparse subset of real images [33], generating synthetic data [10], or relying on simulations based on conventional DEMs [26]. These approaches typically bypass the pushbroom camera model by approximating it with simple pinhole cameras or Rational Polynomial Coefficients (RPCs), failing to faithfully capture the true lunar 3D geometry. Neural volumetric rendering (*e.g.*, NeRF-like approaches) has shown great success in capturing high-fidelity scene geometry, but has not been successfully adapted to these lunar-specific challenges. To properly leverage this rendering paradigm and close this gap, we propose the **Lunar Neural Elevation Model (LNEM)**¹, which enables geometrically consistent lunar surface reconstruction from multi-orbit pushbroom imagery, as illustrated in Fig. 1. Unlike standard neural rendering pipelines, which assume static pinhole cameras, our framework integrates a rigorous sensor model (RSM) that accounts for the pushbroom capture process, per-line camera poses, and changing illumination. This design enables accurate volumetric rendering of the lunar surface by marching each camera ray through a learnable 3D volume, overcoming the common limitations and ambiguities of traditional stereo matching. To validate our approach, we introduce a new multi-orbit lunar dataset that includes precise camera models and linewise rotation matrices. We further present **Lunar Studio**, an RSM-based processing pipeline

¹Our approach to DEM reconstruction differs fundamentally from conventional learning-based DEM refinement methods [20, 30], which train a single network to generalize across large image collections in an image-to-depth paradigm where geometric constraints become weak at test time. Instead, we optimize a volumetric implicit representation that encodes the DEM as the weights of a continuous coordinate-to-density function. This continuous neural field enables nonlinear fusion of heterogeneous orbital observations within a unified coordinate frame while remaining decoupled from image-domain appearance, providing robustness to illumination and albedo variations. In addition, this implicit formulation offers a compact and memory-efficient alternative to large explicit elevation grids.

Table 1. **Key specifications of lunar imaging and altimetry instruments.** Instruments compared are LROC NAC, KPLO LUTI, SELENE TC, and LRO LOLA.

Instrument	Type	Spatial sampling	FOV	Resolution	Spectral range
LROC NAC [27]	dual pushbroom	0.5 m/px (50 km)	2.85°	5064×1	400 to 760 nm
KPLO LUTI [16]	dual pushbroom	2.5 m/px (100 km)	2.88°	2048×1	450 to 850 nm
SELENE TC [13]	stereo pushbroom	10 m/px (100 km)	19.3°	4096×1	430 to 850 nm
LRO LOLA [28]	laser altimeter	5 m spot, 25 m spacing	400 μ rad	–	1064±0.1 nm

Table 2. **Comparison of existing lunar DEMs derived from different instruments.** LOLA provides high vertical accuracy but coarse spatial resolution, while image-based DEMs offer higher resolution at reduced accuracy, with NAC DTMs limited to selected local regions. Limitations are in **bold**. All elevations are relative to the mean lunar radius of 1,737.4 km.

Specification	LDEM [28]	NAC DTM [32]	SLDEM [2]
Instruments	LRO LOLA	LROC NAC	SELENE TC + LRO LOLA
Pixel scale (m/px)	118	2 to 5	60
Vert. accuracy (m)	~1.0	~1	3 to 4
Vert. precision (m)	~0.1	~1	~0.12
Horiz. accuracy (m)	~20	~50	~10
Coverage	global	localized	60°S–60°N

that aligns raw satellite cube images with existing DEMs to produce geometrically consistent inputs for neural rendering. Our contributions are as follows.

- ◊ We introduce the **Lunar Studio** multi-orbit dataset, integrating LROC NAC and KPLO LUTI imagery with LOLA data for evaluation, with automated preprocessing and standardized alignment to mitigate operational cost and distributional bias.
- ◊ We propose **LNEM**, the first volumetric neural rendering framework that integrates Rigorous Sensor Models (RSMs) for learning-based lunar DEM reconstruction, unifying physical sensor geometry with differentiable rendering for heterogeneous pushbroom imagery.
- ◊ Extensive experiments demonstrate that LNEM achieves geometrically consistent reconstruction and consistent performance across different sensors under varying illumination, providing a scalable alternative to conventional DEM reconstruction.

2. Related Work

Lunar digital elevation models and instruments Digital elevation models (DEMs) [2, 28, 32] are foundational for lunar surface analysis, enabling scientific investigation, landing site selection, and mission planning. We summarize instrument specifications in Table 1 and representative DEM products in Table 2.

LOLA [28, 29] provides highly reliable laser altimetry for geometric calibration, integrated into the near-global LDEM. LROC NAC [5, 27] captures dual-pushbroom imagery at 0.5 to 2 m/px for high-resolution local topography, though NAC DTMs cover only selected regions. SELENE TC [13] acquires stereo pushbroom imagery at

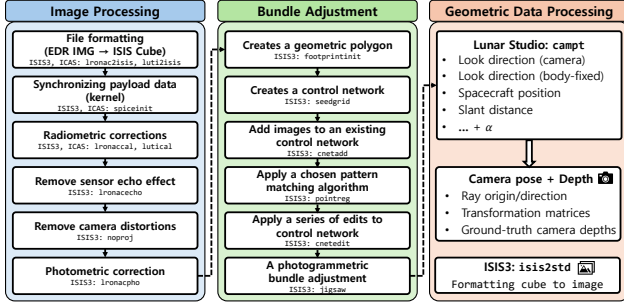


Figure 2. **Overview of Lunar Studio data processing pipeline.** An end-to-end pipeline integrating photometric corrections and rigorous sensor modeling for multi-orbit image alignment, exporting per-line camera rotations and depth products for downstream neural rendering.

10 m/px and, fused with LOLA, produces SLDEM2015 [2] at $\sim 60 m/px$ with improved geometric consistency.

Lunar datasets and 3D vision pipelines Prior neural reconstruction methods for lunar terrain [19, 25, 33] rely on sparse real images, synthetic data, or simulator-derived datasets, bypassing rigorous pushbroom geometry. Processing tools such as ISIS3 [17] and ASP [4] exist but expose fragmented workflows that lack standardized outputs for learning-based pipelines. We address this gap with **Lunar Studio**, an end-to-end ISIS3-based pipeline that exports per-line camera rotations, photometrically corrected imagery for multi-orbit neural rendering.

Neural radiance fields for satellite imagery Neural radiance fields (NeRFs) [23] represent continuous volumetric density and radiance, enabling high-fidelity 3D reconstruction from multi-view images. However, standard NeRFs assume a static pinhole model and are not directly applicable to pushbroom cameras, where each scan line has a distinct projection center [38], violating conventional epipolar geometry [6, 11]. Recent satellite NeRF methods [7, 21, 22] address Earth-observation data by incorporating RPC camera models and illumination modeling, but RPC approximations introduce geometric errors that compound in sparse-view. LunarNRM [33] applies NeRF-style models to the Moon but relies on RPC cameras and lacks a rigorously validated multi-orbit pushbroom benchmark. Other lunar volumetric studies [19, 25] use pinhole-style imagery rather than real orbital pushbroom observations. LNEM addresses these limitations by integrating per-line rigorous sensor models directly into the volumetric rendering pipeline.

3. Lunar Studio and Multi-Orbit Dataset

Motivation on workflow integration We present Lunar Studio, a unified pipeline that integrates fragmented

ISIS3 [17] and ASP [4] workflows for systematic photometric and geometric preprocessing of lunar multi-orbit datasets. Although these tools offer powerful capabilities, their fragmented command structures hinder consistent large-scale processing and lack explicit outputs of key intermediate products such as camera rotation matrices. Lunar Studio resolves this by consolidating the required ISIS3 operations and exporting corrected camera rotations for direct use in downstream learning-based reconstruction.

Motivation on rigorous sensor modeling Rigorous sensor model (RSM) is crucial for accurate 3D reconstruction of planetary surfaces. Pushbroom sensors introduce linewise projection-center shifts, leading to parallax and geometric inconsistencies, while lunar regolith exhibits reflectance variations driven by phase angle, solar incidence, and material composition [14, 18, 38]. Lunar Studio incorporates RSM throughout the pipeline to explicitly model these sensor and illumination effects. To our knowledge, it is the first dedicated pipeline for lunar neural rendering that integrates rigorous sensor modeling, bridging planetary remote sensing and learning-based terrain reconstruction.

3.1. Lunar Studio

Pipeline Our pipeline, illustrated in Fig. 2, integrates key ISIS3 operations to produce consistent, well-aligned inputs for multi-view neural rendering from heterogeneous lunar datasets. To ensure photometric fidelity, we apply radiometric corrections to remove detector bias and dark current and to perform flat-field normalization, followed by photometric correction using an empirical reflectance model built from over 760,000 NAC tiles. Geometric consistency is achieved through multi-stage bundle adjustment: we initialize SPICE [1] geometry with SLDEM2015 as the shape model, refine a multi-image control network, and then run *jigsaw* [8], substantially improving cross-orbit alignment.

3.2. Dataset

Hierarchy Our dataset is organized into a three-level hierarchy: *region*, *site*, and *image*. A *region* represents a broader lunar area of scientific interest. Each *site* is a localized sub-area within a region, selected to capture terrain variability and to collect overlapping observations from multiple orbital passes. An *image* corresponds to a single NAC or LUTI frame that contributes to the site’s multi-view dataset. This hierarchical structure reflects both the spatial organization of lunar terrain and the observational relationships inherent in multi-orbit sensing.

Selected regions We focus on mid-latitude regions that incorporate diverse baseline shape models such as LDEM, SLDEM, and NAC DTM to ensure comprehensive terrain coverage. As summarized in Table 3 and Table 4, each

Table 3. **Terrain characteristics and LROC NAC imagery statistics for the eight benchmark regions in Lunar Studio.** The upper rows list common terrain attributes, and the lower rows summarize NAC image properties. Heights are relative to a mean lunar radius of 1,737,400 *m*. All angular measurements and heights include standard deviations in parentheses.

Region	Apollo 15	Apollo 16	Apollo 17	Eimmart A	Tycho	V. Schröteri	Lacus Mortis Pit	Marius Hills Pit
Terrain type	Mare–Highland	Highlands	Volcanic	Mare Basalt	Impact Crater	Volcanic	Mare Pit Crater	Mare Pit Crater
Lat., Lon.	26.1°N, 3.6°E	9.0°S, 15.7°E	20.2°N, 31.0°E	24.0°N, 65.6°E	43.3°S, 11.4°W	25.5°N, 53.6°W	44.9°N, 25.6°E	14.1°N, 56.7°W
Sites	5	5	5	5	5	5	1	1
Images per site	3	2	3	2	3	3	3	3
LOLA points	22	11	20	109	17	61	21	42
Resolution (<i>px</i>)	512×512	512×512	512×512	512×512	512×512	512×512	512×512	512×512
Pixel scale (<i>m/px</i>)	1.03 to 1.42	0.98 to 1.03	0.98 to 1.45	0.46 to 1.45	0.53 to 0.72	1.21 to 1.49	0.87 to 1.66	0.83 to 1.27
Height (<i>m</i>)	−2013.9 (±107.0)	91.1 (±109.6)	−2335.2 (±626.2)	−2127.3 (±73.2)	−3159.9 (±400.7)	−1649.1 (±57.9)	−2040.9 (±15.0)	−1567.9 (±7.0)
Phase (°)	45.04 (±1.44)	65.66 (±18.85)	64.67 (±5.64)	50.15 (±7.44)	51.38 (±8.91)	49.45 (±7.19)	55.15 (±4.25)	18.90 (±7.46)
Incidence (°)	45.71 (±2.74)	67.61 (±0.99)	61.69 (±6.32)	47.98 (±2.81)	53.15 (±2.54)	52.42 (±0.69)	53.46 (±3.54)	35.36 (±9.77)
North azimuth (°)	243.90 (±2.78)	274.83 (±0.21)	257.45 (±2.40)	246.80 (±1.64)	318.43 (±5.22)	110.83 (±0.56)	139.25 (±10.16)	113.43 (±7.44)

Table 4. **KPLO LUTI imagery statistics in the Lunar Studio Benchmark.** The four regions are processed and summarized.

Region	Apollo 15	Apollo 17	Eimmart A	Tycho
Sites	5	5	5	5
Images per site	3	2	2	2
LOLA points	306	158	98	337
Resolution (<i>px</i>)	512×512	512×512	512×512	512×512
Pixel scale (<i>m/px</i>)	2.62 to 3.68	2.58 to 2.88	2.33 to 2.65	2.18 to 2.62
Height (<i>m</i>)	−1805.4 (±362.3)	−1851.0 (±681.0)	−1885.4 (±173.0)	−2913.1 (±913.0)
Phase (°)	26.72 (±2.78)	20.25 (±0.93)	24.59 (±4.85)	73.51 (±6.39)
Incidence (°)	28.70 (±0.39)	22.50 (±0.30)	42.14 (±0.42)	56.81 (±0.61)
North azimuth (°)	148.34 (±1.50)	147.11 (±1.16)	116.48 (±0.43)	51.12 (±1.00)

region contains multiple NAC and LUTI sites, with each site comprising multi-orbit imagery acquired under varying photometric and geometric conditions. The entire dataset, including images and SLDEM, NAC DTM, and LOLA elevation products, is co-registered within a unified geodetic coordinate frame to ensure spatial consistency.

4. Lunar Neural Elevation Model

Applying neural volumetric rendering [23, 34] to the lunar domain faces three challenges: pushbroom instruments such as LROC NAC [27] provide only grayscale measurements, the lunar surface is dominated by similar craters and low albedo contrast that reduces the semantic and geometric signals needed for reliable reconstruction, and multi-orbit observations yield far fewer viewpoints than typical Earth-observation datasets [9]. We present the LNEM design choices that explicitly address these challenges. LNEM serves as a foundational baseline for pushbroom neural rendering on the Moon, complementing rather than replacing traditional DEM pipelines [2, 3, 15].

Grayscale volume rendering In the volume rendering formulation [23], the color of a pixel is obtained by integrating radiance weighted by density along a camera ray:

$$\hat{C}(\mathbf{r}) = \sum_{i=1}^N T_i \left(1 - \exp(-\sigma_i \delta_i)\right) c_i \quad (1)$$




Figure 3. **LNEM pipeline.** A pushbroom orbit sample is processed to output volume density and a 256-dimensional feature vector. The feature is combined with the sinusoidal-encoded view direction $\gamma(\mathbf{d})$ to predict grayscale c_g , and with the solar direction ω to predict sun visibility v , while I_{ind} is predicted solely from ω . Hidden dimensions are 256 for the shared trunk and 128, 64, and 32 for the subsequent branches, each producing a scalar output.

where the transmittance $T_i = \exp(-\sum_{k<i} \sigma_k \delta_k)$, and σ_i , c_i denote the raw density and intensity at each sampling interval δ_i , respectively. Replacing c_i with a single-band scalar enables the network to learn a grayscale radiance field from monochromatic pixels [35, 37].

Pushbroom ray computation Unlike pinhole cameras where all pixels share a single projection center, pushbroom sensors acquire each line at a distinct time with its own camera position and orientation, and we explicitly model both ray origin and direction in a line-wise manner. For each pushbroom line $l \in \{0, \dots, n\}$, the ephemeris time is $t_l = t_0 + l \Delta t$, where t_0 and t_n denote the start and end times of the image acquisition, and Δt is the line exposure duration. The camera center $\mathbf{o}(l)$ is obtained from the SPK kernel as the spacecraft position at t_l in the Moon body-fixed frame MOON_ME, while the camera-to-Moon rotation $\mathbf{R}_{\text{C2M}}(l)$ is derived from the CK/FK kernels, following the SPICE convention [1], yielding a per-line camera pose.

Given pixel coordinates (sample s , line l), the camera-frame line-of-sight direction $\mathbf{d}(s)$ is computed by converting sample coordinates into focal-plane coordinates using

focal length f , applying distortion corrections from the instrument kernel to obtain corrected coordinates (x_u, y_u) , and normalizing:

$$\mathbf{d}(s) = \frac{(x_u, y_u, -f)}{\sqrt{x_u^2 + y_u^2 + f^2}} \quad (2)$$

The complete ray for pixel (s, l) is then formulated as

$$\mathbf{r}(s, l) = \mathbf{o}(l) + \lambda \mathbf{R}_{\text{C2M}}(l) \mathbf{d}(s) \quad (3)$$

where $\mathbf{d}(s)$ is the camera-frame line-of-sight direction from the RSM, and $\lambda \geq 0$ is the depth.

Positional encoding Standard fixed Fourier feature mappings [31] are inadequate for lunar pushbroom imagery, requiring site-specific high-frequency tuning that leads to slow convergence and inaccurate geometry.

To overcome these limitations, we adopt multi-resolution hash encoding [24], which maps grid-cell corners to shared trainable entries across L levels and concatenates the interpolated per-level features:

$$\mathbf{y}(\mathbf{x}) = (\mathbf{y}_1(\mathbf{x}), \dots, \mathbf{y}_L(\mathbf{x})) \quad (4)$$

which is fed into an MLP. We employ a significantly deeper hierarchy with $L = 32$ levels, twice the default configuration, to capture fine-scale geometric variations, yielding more stable convergence and sharper reconstructions.

Shadow modeling Motivated by shadow-aware irradiance modeling to Earth satellite imagery [7, 21], we perform shadow modeling for lunar terrain. Using the Sun direction derived from Lunar Studio azimuth and incidence angles, we model sun visibility and indirect illumination for each sample along the camera ray via additional branches:

$$c_i = c_{g,i} \cdot (v(\mathbf{x}_i, \boldsymbol{\omega}) + (1 - v(\mathbf{x}_i, \boldsymbol{\omega})) \cdot I_{\text{ind}}(\boldsymbol{\omega})) \quad (5)$$

where $\boldsymbol{\omega}$ is the Sun direction derived from the per-line ephemeris time t_l in Eq. (3), $c_{g,i} \in [0, 1]$ is the base grayscale intensity at sample i , $v(\mathbf{x}_i, \boldsymbol{\omega}) \in [0, 1]$, denoted by v_i , is the sun visibility scalar, and $I_{\text{ind}}(\boldsymbol{\omega}) \in [0, 1]$ is the indirect illumination scalar predicted by a separate MLP conditioned on the Sun direction $\boldsymbol{\omega}$. The full model architecture is illustrated in Fig. 3. We further align solar-ray transmittance $T_{\text{SR},i}$ with the predicted visibility v_i via a shadow correction loss:

$$\mathcal{L}_{\text{SC}}(\mathcal{R}_{\text{SR}}) = \sum_{\mathbf{r} \in \mathcal{R}_{\text{SR}}} \sum_{i=1}^{N_{\text{SR}}} (T_{\text{SR},i} - v_i)^2 \quad (6)$$

For lunar imagery with extreme lighting variations and sparse views, we use only point-wise matching without integral constraints, as the sigmoid-bounded $v_i \in [0, 1]$ provides sufficient regularization. Here, \mathcal{R}_{SR} is the batch of solar rays cast toward the Sun from each sample point along the camera ray, with $N_{\text{SR}} = 512$ samples per solar ray.




Figure 4. **Comparison of 2D and 3D IoU (Intersection over Union).** We quantify spatial overlap between different orbital images. The 2D IoU uses area intersection over union, while the 3D IoU computes volumetric overlap by placing spheres of radius r (ranging from 0 to 10 m) around each 3D point in Moon body-fixed coordinates (MOON_ME) and measuring volumetric overlap across orbits. Eimmart A exhibits lower IoUs due to pixel scale variations from the large orbital altitude difference.




Figure 5. **Comparison before and after bundle adjustment (BA).** Before BA, cross-orbit images exhibit geometric misalignment across corresponding locations, resulting in large ℓ_1 error spikes. The residual errors can be attributed to variations in solar illumination between orbits (solar angular separation $\sim 2.35^\circ$, solar distance difference ~ 0.033 AU).

Depth supervision Dense multi-view constraints are impractical for lunar reconstruction due to sparse orbital passes and large pose variations, making depth supervision an appealing complement to photometric loss. Among available sources, LOLA measurements are too sparse for dense guidance and LDEM is too coarse relative to NAC resolution. NAC DTM provides the highest resolution and quality but covers only selected regions, so we adopt it as the primary depth source where available, falling back to SLDEM which offers the best balance of coverage and resolution elsewhere, as summarized in Table 2.

The ground-truth depth is the distance between camera position $\mathbf{o}(l)$ in Eq. (3) and the DEM surface intersection, with the bounding volume defined as an axis-aligned box from the per-site depth range. Depth supervision also provides metric scale anchoring, preventing the scale ambiguity that arises in photometric-only training. We sample $N = 512$ points per camera ray and optimize:

$$\mathcal{L} = \sum_{\mathbf{r} \in \mathcal{R}} \|\hat{C}(\mathbf{r}) - C(\mathbf{r})\|_2^2 + w_{\text{D}} \|\hat{D}(\mathbf{r}) - D(\mathbf{r})\|_2^2 + w_{\text{SC}} \mathcal{L}_{\text{SC}}(\mathcal{R}_{\text{SR}}) \quad (7)$$

where C and \hat{C} denote the ground-truth and predicted grayscale intensities, D and \hat{D} denote the ground-truth and predicted depths, w_{D} is the depth supervision loss weight, and w_{SC} is the shadow correction loss weight. We set $w_{\text{D}} = 300$ and $w_{\text{SC}} = 0.02$.

Table 5. **Elevation error (m) with respect to LOLA measurements across eight LROC NAC regions.** SLDEM and NAC DTM are production-level DEMs*. Bias-corrected metrics (bias, $\text{RMSE}_{\text{corr}}$, std) are reported for all other methods using 2 to 3 orbits per site. EO-NeRF and Sat-NeRF use RPC models from ASP’s `cam2rpc` on identical Lunar Studio data. SM denotes shadow modeling.

		Apollo 15	Apollo 16	Apollo 17	Eimmart A	Tycho	V. Schröteri	Lacus Mortis Pit	Marius Hills Pit
$\text{RMSE}_{\text{LOLA}}$	SLDEM*	2.115	1.366	1.893	2.842	3.476	5.096	2.487	0.922
	NAC DTM*	1.918	0.616	0.954	3.586	1.551	3.698	4.865	0.823
	LNEM (without SM)	10.602	1.970	8.264	10.995	3.886	10.904	11.214	1.530
	LNEM (with SM)	8.602	4.630	7.318	10.686	2.117	4.248	10.228	0.676
$\text{RMSE}_{\text{corr}}$	ASP [4]: bias	-2.158	-0.638	0.721	-3.547	-1.330	0.955	-205.634	0.949
	$\text{RMSE}_{\text{corr}}$	3.103	0.324	1.986	2.283	0.672	15.083	109.961	1.891
	std	3.061	0.323	1.986	2.248	0.637	15.069	107.120	1.860
	EO-NeRF [22]: bias	311.575	336.557	169.317	423.434	152.403	175.902	493.293	300.569
	$\text{RMSE}_{\text{corr}}$	58.386	7.341	29.824	97.719	46.951	22.826	62.296	38.243
	std	57.156	6.838	29.819	97.710	45.398	22.288	60.874	38.068
	Sat-NeRF [21]: bias	55.888	11.949	42.528	21.369	7.420	-25.642	37.878	-2.547
	$\text{RMSE}_{\text{corr}}$	39.196	15.956	24.169	38.917	11.207	19.137	8.734	6.209
	std	38.886	15.675	24.163	38.649	11.192	18.994	8.587	5.665
	LNEM (with SM): bias	-8.405	-1.375	-7.059	-8.618	-1.782	-1.375	-10.161	0.017
	$\text{RMSE}_{\text{corr}}$	1.565	4.096	1.228	5.674	0.979	3.689	2.025	0.673
	std	1.564	4.079	1.218	5.658	0.975	3.576	2.020	0.665

* SLDEM co-registers 43,200 TC tiles via two-step ICP with GRAIL-based refinement. NAC DTM uses up to 9 stereo pairs aligned to LOLA via geomorphic feature matching.

Table 6. **Elevation error (m) with respect to LOLA measurements across four KPLO LUTI regions.** Eimmart A has no NAC DTM coverage. The elevated bias in LUTI results stems from re-constructed SPICE kernels with larger pointing uncertainties than the precision smithed kernels used for NAC.

		Apollo 15	Apollo 17	Eimmart A	Tycho
$\text{RMSE}_{\text{LOLA}}$	SLDEM*	2.915	3.233	3.034	4.360
	NAC DTM*	1.901	0.898	-	3.191
	LNEM (without SM)	40.980	7.156	6.434	16.931
	LNEM (with SM)	40.102	5.880	2.573	14.434
$\text{RMSE}_{\text{corr}}$	LNEM (with SM): bias	26.965	0.827	1.258	0.019
	$\text{RMSE}_{\text{corr}}$	29.071	5.608	2.261	14.429
	std	29.064	5.410	2.261	13.835

5. Experiments

Training details We implement LNEM in PyTorch using the Adam optimizer with learning rate 5×10^{-4} , $\beta = (0.9, 0.999)$, $\epsilon = 10^{-8}$, a CosineAnnealingLR scheduler with minimum learning rate 5×10^{-6} , and a batch size of 1,024 rays. Each site is trained for 50,000 to 100,000 iterations, requiring approximately 4 to 8 hours on a single NVIDIA RTX 4090 GPU.

LOLA evaluation To assess the absolute vertical accuracy of each DEM, we compare its elevation values against LOLA laser altimetry, following the evaluation convention of [15]. We perform point-to-point matching: each LOLA measurement within the reference image footprint is matched to its nearest DEM grid cell, and the root mean squared error of the elevation residuals is computed:

$$\text{RMSE}_{\text{LOLA}} = \sqrt{\frac{1}{n} \sum_{i=1}^n (z_{\text{LOLA},i} - z_{\text{DEM},i})^2} \quad (8)$$

where n is the number of matched LOLA points. All values in Tables 5, 6, and 8 are reported in meters. Since raw

$\text{RMSE}_{\text{LOLA}}$ can be dominated by a global vertical offset between the DEM and LOLA reference frame, we additionally report bias-corrected metrics for neural methods, where $\text{bias} = \text{median}(z_{\text{LOLA}} - z_{\text{DEM}})$ captures the global vertical offset. The bias-corrected RMSE is:

$$\text{RMSE}_{\text{corr}} = \sqrt{\frac{1}{n} \sum_{i=1}^n (z_{\text{LOLA},i} - z_{\text{DEM},i} - \text{bias})^2} \quad (9)$$

We also report $\text{std} = \text{std}(z_{\text{LOLA}} - z_{\text{DEM}})$, which isolates local terrain shape fidelity independent of systematic bias.

5.1. Quality of Multi-Orbit Dataset

We evaluate the geometric quality of our multi-orbit dataset using two complementary analyses. First, we report Intersection-over-Union (IoU) scores in Fig. 4, which quantify spatial overlap in both 2D image footprints and 3D point distributions across orbits, highlighting the effect of orbital altitude variation on coverage. Second, due to the line-by-line pose variation in pushbroom imagery, accurate camera modeling is essential for multi-orbit consistency; bundle adjustment within Lunar Studio substantially reduces cross-orbit misalignment, as shown in Fig. 5. Together, these results confirm the geometric consistency of our curated multi-orbit dataset.

5.2. Evaluation on LNEM

We evaluate LNEM through both quantitative and qualitative analyses, including elevation error metrics against LOLA altimetry data in Table 5 and Table 6, and qualitative reconstruction quality in Fig. 6 and Fig. 7.

Comparison with neural rendering baselines To contextualize LNEM’s performance, we compare against ASP stereo [4], EO-NeRF [22], and Sat-NeRF [21], where

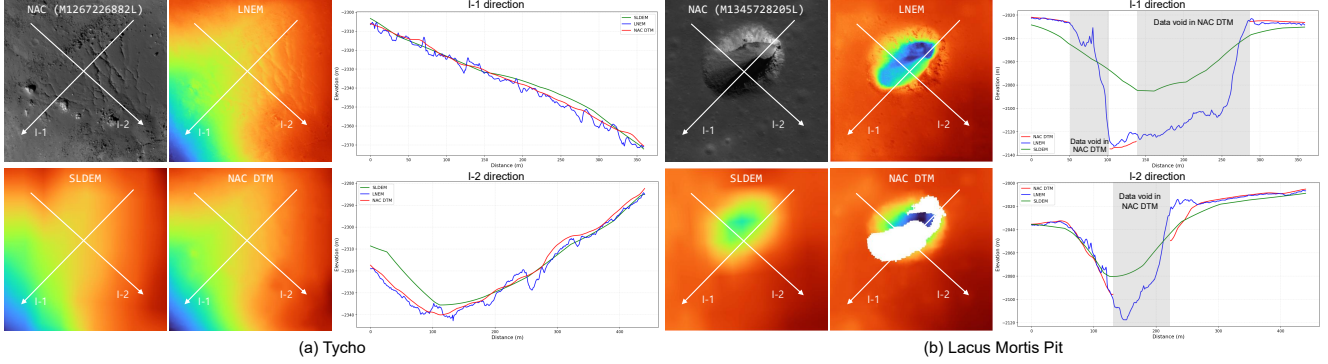


Figure 6. **Qualitative and elevation comparison of SLDEM, NAC DTM, and LNEM, where LNEM is trained using NAC images as input.** We compare DEMs from SLDEM, NAC DTM, and LNEM over (a) Tycho and (b) Lacus Mortis Pit using three-orbit inputs. LNEM reconstructs continuous and detailed elevation, even where NAC DTM exhibits voids. Elevation maps use a common colormap and range, and profiles along I-1 and I-2 are presented.




Figure 7. **Baseline comparison on Lacus Mortis Pit.** (a) Input NAC image. (b) LNEM. (c) Sat-NeRF. (d) EO-NeRF. (e) ASP. Methods (b)–(d) are trained for 50,000 iterations, while (e) is a classical stereo pipeline. Methods (b) and (c) are depth-supervised and share the same supervision data and elevation range, whereas (d) and (e) use separate ranges due to larger offsets. Quantitative results are reported in Table 5.

EO-NeRF and Sat-NeRF use RPC cameras generated via ASP’s `cam2rpc` on identical Lunar Studio data across all eight NAC regions. ASP requires convergence-angle relaxation under small-baseline lunar conditions and still fails at narrow-baseline sites such as Lacus Mortis Pit due to triangulation instability. As shown in Table 5, EO-NeRF relies solely on photometric consistency without depth supervision, resulting in large global vertical offsets due to scale ambiguity, even reconstructing the pit crater as convex. Sat-NeRF mitigates this issue with depth supervision but still exhibits significantly higher $\text{RMSE}_{\text{corr}}$ than LNEM across most regions, highlighting the advantage of rigorous pushbroom sensor modeling over RPC approximations. ASP achieves competitive bias-corrected accuracy in regions with favorable stereo geometry such as Apollo 16 and Tycho, but degrades severely at the challenging sites noted above. Table 5 provides a detailed case study on Lacus Mortis Pit, with qualitative comparison shown in Fig. 7.

Regional variation and bias-corrected analysis Planetary observations inherently vary in quality, coverage, and observation geometry. As shown in Table 5, after bias correction, LNEM achieves $\text{RMSE}_{\text{corr}}$ of 0.67 to 5.67 m across all NAC regions. Sites with three input views achieve 0.67

Table 7. **PSNR comparison between sinusoidal and adaptive hash encodings.** The sinusoidal baseline is sensitive to M , while LNEM remains consistent across sites.

Site (NAC)	Encode	Sinusoidal (Fourier)				Adaptive hash
	$M = 4$	$M = 10$	$M = 16$	$M = 30$	LNEM	
Apollo 15	24.29	28.33	28.25	29.01	48.41	
Apollo 17	26.46	32.92	32.82	33.38	48.39	

to 3.69 m , while two-view sites such as Apollo 16 and Eimmart A show elevated errors of 4.10 to 5.67 m . Eimmart A further exhibits lower 3D IoU due to large orbital altitude variation in Fig. 4, which reduces cross-orbit spatial overlap and degrades geometric constraints. The high raw errors in several regions mainly arise from global vertical offsets, as reflected by the gap between $\text{RMSE}_{\text{LOLA}}$ and $\text{RMSE}_{\text{corr}}$ in Table 5. For KPLO LUTI in Table 6, even after bias correction, Apollo 15 yields $\text{RMSE}_{\text{corr}} = 29.07 m$ and $\text{std} = 29.06 m$, confirming that the error originates from kernel quality rather than the method itself. These results could be further improved with smoothed kernels providing higher pointing precision [36].

5.3. Ablation Studies

Effectiveness of adaptive positional encoding When Fourier frequency M is optimized for one site, the sinusoidal (Fourier) baseline fails to render grayscale intensity correctly on another, confirming its limited generalization and the high cost of site-specific tuning. Adaptive multi-resolution hash encoding in Eq. (4) achieves robust, high-fidelity representations across all sites with a single fixed set of parameters. As shown in Table 7, the best sinusoidal baseline yields inconsistent PSNR across sites depending on M , whereas adaptive multi-resolution hash encoding in Eq. (4) achieves robust, high-fidelity representations across all sites with a single fixed set of parameters.

Table 8. Mean σ_{total}^2 (m^2) across eight regions for LNEM with and without shadow modeling (SM). σ_{total}^2 quantifies the dispersion across orbits in matched 3D points, where lower values indicate higher geometric consistency.

	Apollo 15	Apollo 16	Apollo 17	Eimmart A	Tycho	V. Schröteri	Lacus Mortis Pit	Marius Hills Pit
w/o SM	3.1884	1.1773	5.6028	8.7154	4.2006	4.5983	17.8351	0.7088
w/ SM	2.2145	1.1481	4.3430	7.9491	3.9869	3.5194	6.2285	0.6332




Figure 8. **Geometric consistency** (σ_{total}^2) map (Apollo 15). Shadow modeling reduces multi-view inconsistency, yielding lower σ_{total}^2 values and cleaner deviation maps.

Shadow modeling and geometric consistency To evaluate the impact of shadow modeling on cross-orbit consistency, we analyze both elevation error and geometric consistency. As shown in Table 5, LNEM with shadow modeling achieves lower $\text{RMSE}_{\text{LOLA}}$ than the variant without shadow modeling in 7 of 8 cases. The only exception is Apollo 16, a two-view site where shadow modeling increases $\text{RMSE}_{\text{LOLA}}$ from 1.970 to 4.630 m , suggesting overfitting due to insufficient multi-view constraints. To further quantify the effect of shadow modeling across orbits, we measure geometric consistency by projecting predicted depth maps from multiple viewpoints into the `MOON_ME` frame, yielding one 3D point cloud per orbit. For each reference point, nearest neighbors are retrieved from other orbits within a 50 m threshold. Spatial dispersion is measured using per-axis standard deviations ($\sigma_x, \sigma_y, \sigma_z$), yielding the rotation-invariant metric:

$$\sigma_{\text{total}}^2 = \sigma_x^2 + \sigma_y^2 + \sigma_z^2 \quad (10)$$

Smaller σ_{total}^2 indicates tighter multi-view convergence and higher geometric consistency. As shown in Table 8, shadow modeling improves σ_{total}^2 across all regions. The spatial distribution of σ_{total}^2 in Fig. 8 further reveals that regions with high variance are strongly correlated with severe photometric variation, particularly near shadow boundaries where cross-orbit depth estimation becomes ambiguous.

Sensitivity to depth supervision We vary the percentage of rays receiving depth supervision on Apollo 17. As shown in Fig. 9, all ratios achieve comparable reconstruction quality, confirming that LNEM does not memorize the supervising DEM but learns geometry primarily through multi-view photometric consistency.




Figure 9. **Depth supervision sensitivity** (Apollo 17). Qualitative DEM results at 10%, 40%, and 100% supervision ratios, achieving $\text{RMSE}_{\text{CORR}}$ of 2.444 m , 2.098 m , and 1.864 m , respectively.




Figure 10. **Shadow map prediction** (Eimmart A). Brighter regions indicate areas directly illuminated by the Sun, while non-overlapping regions exhibit unreliable estimates.

6. Conclusion

A key motivation behind our research is to facilitate the development of neural DEM approaches that complement existing methods [2, 3, 15]. Toward this end, we introduce Lunar Studio, a data processing pipeline using RSM-based pushbroom imagery from LROC NAC and KPLO LUTI, and publicly release a curated dataset featuring eight near-side regions of significant geological and scientific interest. Building on this, we propose LNEM, which leverages continuous volumetric representations and multi-view parallax under rigorous sensor models to address the challenges of sparse observations, grayscale imagery, and subtle textural variations, achieving geometrically consistent reconstruction under varying illumination conditions.

Limitations and future work While LNEM demonstrates promising results for lunar DEM reconstruction, our shadow modeling approach relies on photometric intensity-based learning rather than explicitly modeling the physical reflectance properties of lunar regolith. As illustrated in Fig. 10, this approach does not fully capture secondary illumination effects and polarization characteristics that vary across terrains. Incorporating physics-based reflectance models, such as the Hapke BRDF [12], is a promising direction for improving physical fidelity. Currently, our benchmark is focused on mid-latitude regions. Future work will extend the dataset to include south polar NAC observations and enable multi-sensor fusion across heterogeneous sources, such as NAC and LUTI, within a unified model.

Acknowledgements

This work was supported by the Science Research Program using Danuri's science payloads (No. 2026185100) funded by the Korea Astronomy and Space Science Institute (KASI). This research was supported by Basic Science Research Program through the National Research Foundation of Korea (NRF) funded by the Ministry of Education (No. RS-2024-00463470).

References

- [1] Charles H. Acton Jr. Ancillary data services of NASA's navigation and ancillary information facility. *Planetary and Space Science*, 44(1):65–70, 1996. 3, 4
- [2] M. K. Barker, E. Mazarico, G. A. Neumann, M. T. Zuber, Junichi Haruyama, and D. E. Smith. A new lunar digital elevation model from the lunar orbiter laser altimeter and SELENE terrain camera. *Icarus*, 273:346–355, 2016. 2, 3, 4, 8
- [3] Michael K Barker, Erwan Mazarico, Gregory A Neumann, David E Smith, Maria T Zuber, and James W Head. Improved LOLA elevation maps for south pole landing sites: Error estimates and their impact on illumination conditions. *Planetary and Space Science*, 203:105119, 2021. 4, 8
- [4] Ross A. Beyer, Oleg Alexandrov, and Scott McMichael. The ames stereo pipeline: NASA's open source software for deriving and processing terrain data. *Earth and Space Science*, 5(9):537–548, 2018. 3, 6
- [5] Gordon Chin, Scott Brylow, Marc Foote, James Garvin, Justin Kasper, John Keller, Maxim Litvak, Igor Mitrofanov, David Paige, Keith Raney, et al. Lunar reconnaissance orbiter overview: The instrument suite and mission. *Space Science Reviews*, 129:391–419, 2007. 2
- [6] Carlo De Franchis, Enric Meinhardt-Llopis, Julien Michel, Jean-Michel Morel, and Gabriele Facciolo. An automatic and modular stereo pipeline for pushbroom images. In *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 2014. 2, 3
- [7] Dawa Derksen and Dario Izzo. Shadow neural radiance fields for multi-view satellite photogrammetry. In *CVPR Workshops*, 2021. 3, 5
- [8] K. L. Edmundson, D. A. Cook, O. H. Thomas, B. A. Archinal, and R. L. Kirk. Jigsaw: The ISIS3 bundle adjustment for extraterrestrial photogrammetry. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 1:203–208, 2012. 3
- [9] Jian Gao, Jin Liu, and Shunping Ji. Rational polynomial camera model warping for deep learning based satellite multi-view stereo matching. In *ICCV*, 2021. 4
- [10] Clémentine Grethen, Simone Gasparini, Géraldine Morin, Jérémy Lebreton, Lucas Marti, and Manuel Sanchez-Gestido. Adapting stereo vision from objects to 3d lunar surface reconstruction with the stereolunar dataset. In *ICCV Workshops*, 2025. 2
- [11] Rajiv Gupta and Richard I. Hartley. Linear pushbroom cameras. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(9):963–975, 1997. 2, 3
- [12] Bruce Hapke. Bidirectional reflectance spectroscopy: 4. the extinction coefficient and the opposition effect. *Icarus*, 67(2):264–280, 1986. 8
- [13] J. Haruyama, M. Ohtake, T. Matsunaga, T. Morota, Y. Yokota, C. Honda, N. Hirata, H. Demura, A. Iwasaki, R. Nakamura, et al. Planned radiometrically calibrated and geometrically corrected products of lunar high-resolution terrain camera on SELENE. *Advances in Space Research*, 42(2):310–316, 2008. 2
- [14] Paul Helfenstein and Joseph Veverka. Photometric properties of lunar terrains derived from hapke's equation. *Icarus*, 72(2):342–357, 1987. 3
- [15] M. R. Henriksen, M. R. Manheim, K. N. Burns, P. Seymour, E. J. Speyerer, A. Deran, A. K. Boyd, E. Howington-Kraus, Mark R. Rosiek, Brent A. Archinal, et al. Extracting accurate and precise topography from LROC narrow angle camera stereo observations. *Icarus*, 283:122–137, 2017. 4, 6, 8
- [16] Moon-Jin Jeon, Young-Ho Cho, Eunhyeuk Kim, Dong-Gyu Kim, Young-Joo Song, SeungBum Hong, Jonghee Bae, Jun Bang, Jo Ryeong Yim, and Dae-Kwan Kim. Korea pathfinder lunar orbiter (KPLLO) operation: From design to initial results. *Journal of Astronomy and Space Sciences*, 41(1):43–60, 2024. 2
- [17] Kelvin Rodriguez and Astrogeology Science Center. Integrated software for imagers and spectrometers (ISIS) 8.3.0, 2024. 3
- [18] M. A. Kreslavsky, Yu. G. Shkuratov, Yu. I. Velikodsky, V. G. Kaydash, D. G. Stankevich, and C. M. Pieters. Photometric properties of the lunar surface derived from clementine observations. *Journal of Geophysical Research: Planets*, 105(E8):20281–20295, 2000. 3
- [19] Suwan Lee and Seokju Lee. Exploring tycho crater: 3d reconstruction with neural radiance fields from sparse views. In *55th Lunar and Planetary Science Conference*, 2024. 3
- [20] Yang Liu, Yexin Wang, Kaichang Di, Man Peng, Wenhui Wan, and Zhaoqin Liu. A generative adversarial network for pixel-scale lunar dem generation from high-resolution monocular imagery and low-resolution dem. *Remote Sensing*, 14(21):5420, 2022. 2
- [21] Roger Marí, Gabriele Facciolo, and Thibaud Ehret. Sat-NeRF: Learning multi-view satellite photogrammetry with transient objects and shadow modeling using RPC cameras. In *CVPR Workshops*, 2022. 3, 5, 6
- [22] Roger Marí, Gabriele Facciolo, and Thibaud Ehret. Multi-date earth observation NeRF: The detail is in the shadows. In *CVPR Workshops*, 2023. 3, 6
- [23] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. 3, 4
- [24] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.*, 41(4):102:1–102:15, 2022. 5
- [25] Art Prosvetov, Alexander Govorov, Maxim Pupkov, Alexander Andreev, and Vladimir Nazarov. Illuminating the moon:

- Reconstruction of lunar terrain using photogrammetry, neural radiance fields, and gaussian splatting. *Astronomy and Computing*, 52:100953, 2025. 3
- [26] Antoine Richard, Junnosuke Kamohara, Kentaro Uno, Shreya Santra, Dave Van Der Meer, Miguel Olivares-Mendez, and Kazuya Yoshida. Omnirls: A photorealistic simulator for lunar robotics. In *ICRA*, 2024. 2
- [27] Mark S. Robinson, S. M. Brylow, M. Tschimmel, D. Humm, S. J. Lawrence, P. C. Thomas, Brett W. Denevi, E. Bowman-Cisneros, J. Zerr, M. A. Ravine, et al. Lunar reconnaissance orbiter camera (LROC) instrument overview. *Space Science Reviews*, 150:81–124, 2010. 2, 4
- [28] David E Smith, Maria T Zuber, Glenn B Jackson, John F Cavanaugh, Gregory A Neumann, Haris Riris, Xiaoli Sun, Ronald S Zellar, Craig Coltharp, Joseph Connelly, et al. The lunar orbiter laser altimeter investigation on the lunar reconnaissance orbiter mission. *Space Science Reviews*, 150:209–241, 2010. 2
- [29] David E Smith, Maria T Zuber, Gregory A Neumann, Frank G Lemoine, Erwan Mazarico, Mark H Torrence, Jan F McGarry, David D Rowlands, James W Head III, Thomas H Duxbury, et al. Initial observations from the lunar orbiter laser altimeter (LOLA). *Geophysical Research Letters*, 37(18), 2010. 2
- [30] Corinne Stucker and Konrad Schindler. Resdepth: A deep residual prior for 3d reconstruction from high-resolution satellite images. *ISPRS Journal of Photogrammetry and Remote Sensing*, 183:560–580, 2022. 2
- [31] Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *NeurIPS*, 2020. 5
- [32] T. Tran, M. R. Rosiek, Ross A. Beyer, S. Mattson, E. Howington-Kraus, M. S. Robinson, B. A. Archinal, K. Edmundson, D. Harbour, E. Anderson, et al. Generating digital terrain models using LROC NAC images. In *ASPRS/CaGIS 2010 Fall Specialty Conference*, 2010. 2
- [33] Ellemieke Van Kints, Aiden Hammond, Caleb Adams, and Ignacio G Lopez-Francos. Neural radiance methods for lunar terrain modeling. In *Aerospace Conference*, 2025. 2, 3
- [34] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. NeuS: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *NeurIPS*, 2021. 4
- [35] Xin Wang, Yineng Chen, Shu Hu, Heng Fan, Hongtu Zhu, and Xin Li. Neural radiance fields in medical imaging: A survey. *arXiv preprint arXiv:2402.17797*, 2024. 4
- [36] Bo Wu and Wai Chung Liu. Calibration of boresight offset of lroc nac imagery for precision lunar topographic mapping. *ISPRS Journal of Photogrammetry and Remote Sensing*, 128:372–387, 2017. 7
- [37] Tianxiang Ye, Qi Wu, Junyuan Deng, Guoqing Liu, Liu Liu, Songpengcheng Xia, Liang Pang, Wenxian Yu, and Ling Pei. Thermal-NeRF: Neural radiance fields from an infrared camera. In *IROS*, 2024. 4
- [38] Jingyi Yu, Leonard McMillan, and Peter Sturm. Multiperspective modeling, rendering, and imaging. In *ACM SIGGRAPH ASIA 2008 Courses*, pages 1–36. Association for Computing Machinery, 2008. 3