

LNEM: Lunar Neural Elevation Model

Supplementary Material

A. Overview of Supplementary Material

This supplementary material provides additional implementation details, ablation studies, and analyses to support the main paper. All tables, figures, equations, and references in this supplementary file are self-contained. The sections are organized as follows:

- **Sec. B:** Geometric data processing efficiency comparison between ISIS3 and Lunar Studio.
- **Sec. C:** Bundle adjustment statistics across all eight regions.
- **Sec. D:** Comparison between fixed Fourier features and multi-resolution hash encoding.
- **Sec. E:** Shadow modeling analysis at the Eimmart A and Apollo 15 sites.
- **Sec. F:** Full $\text{RMSE}_{\text{LOLA}}$ evaluation across all regions, including a Lacus Mortis Pit case study.
- **Sec. G:** Additional elevation profile analyses and depth supervision sensitivity.
- **Sec. H:** Site selection criteria and regional characteristics of the Lunar Studio benchmark.
- **Sec. I:** LUTI LNEM performance analysis and the impact of SPICE kernel quality.
- **Sec. J:** Limitations and future directions.

B. ISIS3 and Lunar Studio Comparison

To enable efficient large-scale dataset preparation for LNEM, which requires processing millions of surface points, we developed *Lunar Studio*, a custom geometric data processing tool that extends the functionality of ISIS3 `campt`. A comparison between the two approaches is summarized in Table A1.

Standard ISIS3 `campt` requires pre-defined coordinate text files and computes all geometric parameters for each query, regardless of whether they are needed, resulting in unnecessary computational overhead. In contrast, Lunar Studio introduces two key improvements: (1) automatic parameter specification and (2) on-demand computation of se-

lected geometric attributes. Rather than exhaustively computing all outputs, our method selectively evaluates only the required parameters, such as incidence, emission, and phase angles. This on-demand strategy significantly reduces redundant computation and achieves a 10–20% reduction in preprocessing time, as summarized in Table A1. This efficiency gain is particularly important when scaling to millions of surface samples in the LNEM dataset, as detailed in Table 3 and Table 4 of the main paper.

C. Bundle Adjustment Statistics

As reported in Table A2, *jigsaw* bundle adjustment achieves sub-pixel root mean square (RMS) residuals across most LROC NAC regions. Bundle adjustment primarily removes large-scale geometric bias in camera poses, enabling consistent multi-view geometry.

For Apollo 17, $\text{RMSE}_{\text{LOLA}}$ decreases from 73.4 to 7.3 *m* after bundle adjustment, while the local terrain fidelity std improves from 2.53 to 1.22 *m*, demonstrating substantial improvement in both global alignment and local shape accuracy. Marius Hills Pit exhibits the highest bundle adjustment RMS of approximately 4.1 pixels, yet achieves the lowest elevation error in Table 5 of the main paper. This indicates that residual pixel-level misalignment does not necessarily degrade local geometric reconstruction, highlighting the robustness of the continuous latent representation.

The LUTI Apollo 15 residual of 1.38 pixels is 14 times larger than that of NAC Apollo 15 at 0.098 pixels. This directly explains the elevated $\text{RMSE}_{\text{LOLA}}$ of 40.1 *m* for LUTI Apollo 15 compared to 8.6 *m* for NAC Apollo 15, confirming that SPICE kernel quality fundamentally bounds reconstruction accuracy in our ray-based framework.

D. Fixed Fourier Features and Multi-Resolution Hash Encoding Comparison

We analyze the impact of input coordinate encoding on reconstructing high-frequency geometric details from single-

Table A1. **Geometric data processing efficiency: ISIS3 `campt` vs. Lunar Studio.** Lunar Studio enables on-demand computation with automatic parameter generation, reducing preprocessing time by 10 to 20% compared to the standard `campt` pipeline.

Processing tool	Processing unit	Operation strategy	Parallelization	Specification overhead
ISIS3: <code>campt</code> (previous approach)	specified pixel, all geometric data	batch processing (full parameters)	supported	manual definition (via coordinate txt file)
Lunar Studio (ours, customized)	specified pixel, selected geometric data	on-demand (selective parameters)	supported	automatic generation (via command parameters)

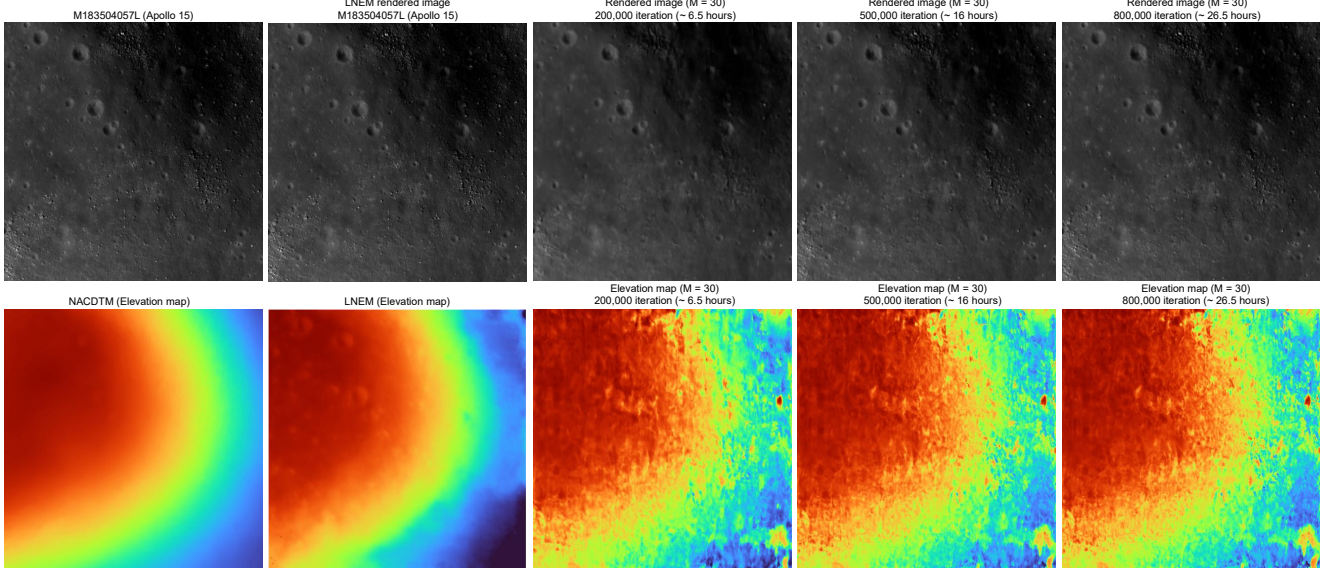


Figure A1. **Qualitative results of Fourier feature mapping across different training iterations at the Apollo 15 site.** With frequency hyperparameter $M = 30$, the rendered grayscale imagery converges to a blurry state and the elevation map fails to capture valid geometry. This confirms that failure is not due to insufficient training time but rather a representational bottleneck in capturing the high-frequency components of lunar terrain. These results demonstrate that adaptive encoding is more suitable for lunar surface reconstruction than standard Fourier feature mapping.

Table A2. **Bundle adjustment RMS residuals (px) across eight LROC NAC regions.** Residuals are reprojection errors of tie points, where lower values indicate better geometric alignment.

	Apollo 15	Apollo 16	Apollo 17	Eimmart A	Tycho	V. Schröteri	Lacus Mortis Pit	Marius Hills Pit
BA RMS	0.098	0.142	0.099	0.333	1.412	0.796	1.337	4.143

channel lunar imagery. Standard NeRF methods employ fixed Fourier feature mapping to mitigate spectral bias [11]. However, this deterministic encoding relies on a fixed set of frequency bands, which limits its ability to adapt to the spatially varying complexity of lunar terrain. Lunar surface imagery contains diverse geometric and radiometric structures [5], ranging from smooth mare regions to highly detailed crater boundaries. A fixed-frequency representation struggles to simultaneously capture both low- and high-frequency variations across such heterogeneous regions.

The traditional positional encoding $\gamma(\mathbf{p})$ maps each coordinate of the input position vector \mathbf{p} to a higher-dimensional space using a predefined set of sinusoidal functions:

$$\gamma(p) = \left(\sin(2^0 \pi p), \cos(2^0 \pi p), \dots, \sin(2^{M-1} \pi p), \cos(2^{M-1} \pi p) \right) \quad (\text{A1})$$

where p denotes a scalar coordinate of the input position \mathbf{p} , and the encoding is applied independently to each coordinate. The hyperparameter M controls the maximum fre-

Table A3. **Hash encoding parameters.** A fixed parameter set consistently produces high-quality reconstructions across all tested sites without any site-specific adjustment.

Parameter	Symbol	Value
Number of levels	L	32
Max. entries per level (hash table size)	T	2^{23}
Number of feature dimensions per entry	F	2
Coarsest resolution	N_{\min}	16
Finest resolution	N_{\max}	524288

quency band. For reference, the same encoding is applied to the view direction \mathbf{d} with $M = 4$ in Fig. 3 of the main paper. A key limitation of this formulation is that the frequency basis is fixed across all spatial locations, requires tuning of M for each target region, and fails to generalize across scenes with different geometric complexity.

Efficacy of adaptive hash encoding To overcome the limitations of fixed Fourier feature mappings, LNEM leverages multi-resolution hash encoding [9]. Unlike fixed-frequency embeddings, this adaptive representation accommodates spatially varying geometric complexity without requiring site-specific hyperparameter tuning. We employ a single parameter set in Table A3 across all sites, achieving consistent performance across diverse lunar datasets.

As shown in Fig. A1 and Fig. A6, the Fourier-based

Table A4. $\text{RMSE}_{\text{LOLA}} (m)$ with respect to LOLA measurements across eight LROC NAC regions. $\text{RMSE}_{\text{LOLA}}$ reflects both systematic vertical offsets (bias) and local shape errors, while bias-corrected metrics ($\text{RMSE}_{\text{corr}}$, std) provided in Table 5 of the main paper isolate local geometric accuracy.

	Apollo 15	Apollo 16	Apollo 17	Eimmart A	Tycho	V. Schröteri	Lacus Mortis	Marius Hills
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
ASP [1]	4.058	0.693	2.127	3.869	1.287	15.154	210.150	2.262
EO-NeRF [7]	305.771	357.258	171.402	433.389	148.309	183.152	483.904	299.401
Sat-NeRF [6]	64.259	19.037	49.365	46.551	13.176	33.959	37.284	5.666
LNEM (w/o SM)	10.602	1.970	8.264	10.995	3.886	10.904	11.214	1.530
LNEM (w/ SM)	8.602	4.630	7.318	10.686	2.117	4.248	10.228	0.676

*SLDEM co-registers 43,200 TC tiles via two-step ICP with GRAIL-based refinement and NAC DTM uses up to 9 stereo pairs aligned to LOLA via geomorphic feature matching. All other methods use 2 to 3 orbits per site.




Figure A2. **Qualitative results of hash encoding resolution levels.** All results are shown at 50k iterations with identical hyperparameters, varying only the number of levels L .

baseline fails to disentangle texture from geometry due to its fixed frequency basis. Crucially, this limitation arises from a fundamental representational deficit rather than insufficient training. Even when extending training to 8×10^5 iterations, the baseline yields negligible improvement in recovering fine geometric details. In contrast, LNEM reconstructs high-frequency geometric and photometric structures within 50,000 iterations, achieving superior data efficiency and robustness without frequency tuning.

Effect of hash encoding resolution levels We use $L = 32$ levels, twice the default configuration, to better capture fine-scale geometric variations in lunar imagery. As shown in Fig. A2, increasing L enhances the representation of high-frequency spatial details, leading to more stable convergence and sharper reconstructions.

E. Shadow Modeling Analysis

Shadow map prediction The predicted shadow map at the Eimmart A site is shown in Fig. 10 of the main paper. Areas with direct solar illumination appear brighter, while regions outside the shared footprint of all orbits yield unreliable visibility estimates due to limited multi-view observations. This behavior is consistent with the shadow model’s tendency to overfit under sparse-view conditions, such as 2-view configurations at Apollo 16 and Eimmart A, as discussed in the main paper.

Table A5. **Baseline comparison on Lacus Mortis Pit.** $\text{RMSE}_{\text{corr}}$ and std are computed after point-to-point LOLA matching with median bias removal.

Method	Valid LOLA pts	$\text{RMSE}_{\text{LOLA}} (m)$	bias (m)	$\text{RMSE}_{\text{corr}} (m)$	std (m)
ASP [1]	6	210.150	-205.634	109.961	107.120
EO-NeRF [7]	21	483.904	493.293	62.296	60.874
Sat-NeRF [6]	21	37.284	37.878	8.734	8.587
LNEM (ours)	21	10.228	-10.161	2.025	2.020

Spatial uncertainty and shadow boundaries Per-pixel σ_{total}^2 maps for LNEM with and without shadow modeling at the Apollo 15 site are shown in Fig. 8 of the main paper. Shadow modeling reduces σ_{total}^2 across the site. The mean σ_{total}^2 decreases from 3.1884 without shadow modeling to 2.2145 with shadow modeling, as reported in Table 8 of the main paper. Regions with the highest residual inconsistency coincide with shadow boundaries, where large photometric variation across orbits makes multi-view depth estimation inherently ambiguous. This spatial pattern confirms that σ_{total}^2 provides a physically meaningful per-pixel uncertainty estimate without requiring additional network components.

F. Baseline Comparison

Table A4 reports $\text{RMSE}_{\text{LOLA}}$ for all methods across all eight NAC regions. $\text{RMSE}_{\text{LOLA}}$ is dominated by global vertical offsets between each method and the LOLA reference frame, making it less suitable for direct method comparison. We additionally provide bias-corrected metrics (bias, $\text{RMSE}_{\text{corr}}$, std) in Table 5 of the main paper to better reflect local geometric accuracy.

EO-NeRF achieves $\text{RMSE}_{\text{LOLA}}$ ranging from 148 to 484 m and Sat-NeRF from 5 to 64 m , reflecting large global offsets arising from the absence of or inaccurate metric depth anchoring rather than poor local geometry. ASP achieves competitive $\text{RMSE}_{\text{LOLA}}$ in geometrically favorable regions, reaching 0.69 m at Apollo 16 and 1.29 m

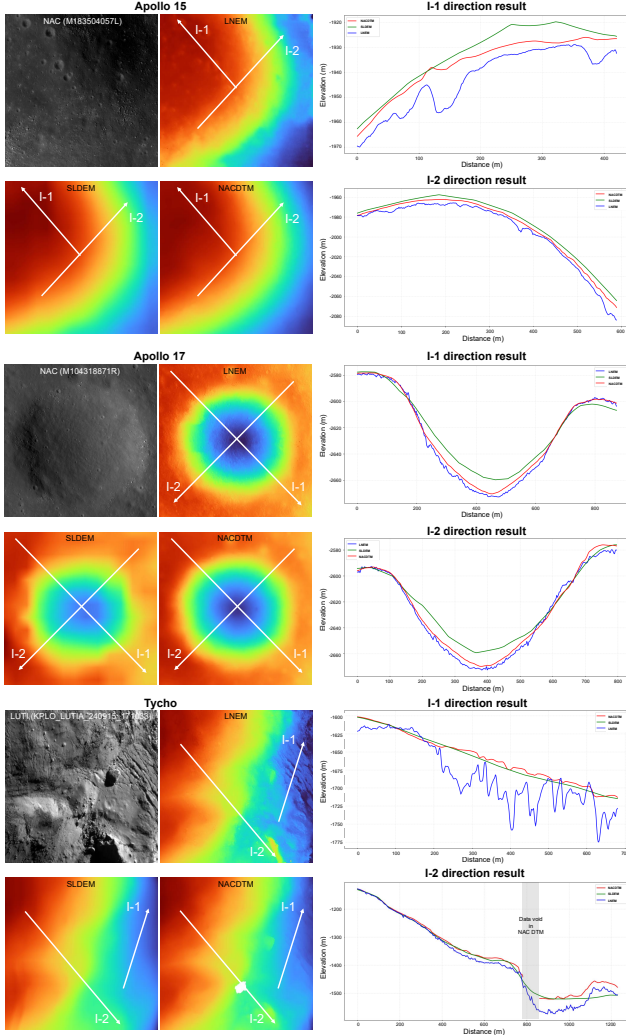


Figure A3. **Qualitative and cross-section elevation comparisons of SLDEM, NAC DTM, and LNEM.** Elevation maps use a common colormap and range, and profiles along I-1 and I-2 are presented, following the same procedure as Fig. 6 in the main paper.

at Tycho. However, its performance degrades at narrow-baseline sites such as Lacus Mortis Pit, where $\text{RMSE}_{\text{LOLA}}$ reaches 210 m due to triangulation failure. LNEM with shadow modeling achieves the lowest $\text{RMSE}_{\text{LOLA}}$ among neural methods, ranging from 0.676 to 10.686 m across NAC regions. For sites where bias correction yields minimal improvement, the residual error is dominated by local shape inaccuracies rather than global vertical misalignment.

Lacus Mortis Pit Table A5 provides detailed results for the Lacus Mortis Pit site, where narrow-baseline geometry and abrupt depth discontinuities within the pit interior make reconstruction particularly challenging. The qualitative elevation comparison is shown in Fig. 7 of the main paper.

ASP stereo requires convergence-angle relaxation under




Figure A4. **Convergence behavior of LOLA mean error.** The right panel shows 22 LOLA measurement points (red dots) at the Apollo 15 site in M183504057L. The left plot illustrates the evolution of the mean error over training iterations. The initial increase in LOLA error for LNEM with shadow modeling indicates that the model first prioritizes learning geometric structure through multi-view photometric consistency rather than fitting the depth prior.

small-baseline lunar conditions and produces only 6 valid LOLA match points due to triangulation failure over the pit interior, with $\text{RMSE}_{\text{LOLA}}$ of 210 m dominated by large systematic bias. After bias correction, std of 107.120 m indicates poor local shape recovery. EO-NeRF, relying solely on photometric consistency without depth supervision, exhibits large errors with bias exceeding 490 m , and even produces inverted elevation structures within the pit.

Sat-NeRF with depth supervision achieves $\text{RMSE}_{\text{corr}}$ of 8.734 m , but its std of 8.587 m is more than four times that of LNEM, highlighting the geometric advantage of rigorous pushbroom sensor modeling over RPC approximations. LNEM achieves the best performance across all metrics, with $\text{RMSE}_{\text{corr}}$ of 2.025 m and std of 2.020 m .

Furthermore, the multi-view geometric consistency metric σ_{total}^2 , defined in Eq. 10 of the main paper, confirms that LNEM achieves substantially tighter cross-orbit agreement with a mean value of 6.229 m^2 , compared to 12.127 m^2 for Sat-NeRF and 12.122 m^2 for EO-NeRF, demonstrating the geometric advantage of RSM-based modeling over RPC approximations.

G. Additional Elevation Profile Analyses

We compare DEMs generated by LNEM with SLDEM and NAC DTM across three sites: Apollo 15 (NAC 2-view), Apollo 17 (NAC 3-view), and Tycho (LUTI 2-view). As shown in Fig. A3, LNEM closely follows the global elevation trends of both baselines. In the I-1 profiles of the Apollo 15 and Tycho regions, LNEM additionally recovers geometric details that are over-smoothed in SLDEM and NAC DTM.

Fig. A4 shows the convergence behavior of the LOLA mean error, computed as the average absolute elevation difference between the reconstructed DEM and corresponding LOLA measurements at matched locations. During early training, LNEM with shadow modeling exhibits a tempo-




Figure A5. **Full depth supervision sensitivity results (Apollo 17).** Qualitative DEM reconstructions are shown for supervision ratios from 10 % to 100 %. The input NAC image (M104318871R) is shown in the top-left panel. All ratios produce visually comparable reconstructions with no systematic degradation, and $\text{RMSE}_{\text{corr}}$ shows no monotonic trend in Table A6.

Table A6. **Depth supervision sensitivity in Apollo 17.** $\text{RMSE}_{\text{corr}}$ (m) at each supervision ratio from 10% to 100% of supervised rays. All ratios achieve comparable accuracy with no monotonic trend, confirming low sensitivity to the supervising DEM.

Supervised rays (%)	10	20	30	40	60	80	100
$\text{RMSE}_{\text{corr}}$ (m)	2.444	2.035	1.614	2.098	2.066	1.737	1.864

rary increase in error, indicating that the model first prioritizes learning geometric structure through multi-view photometric consistency before depth supervision dominates.

Fig. A5 and Table A6 present qualitative DEM results and $\text{RMSE}_{\text{corr}}$ across supervision ratios from 10 % to 100 %. All ratios achieve comparable reconstruction quality with no monotonic trend, with the best performance at 30 % ($\text{RMSE}_{\text{corr}} = 1.614 m$). This confirms that LNEM learns geometry primarily through multi-view photometric consistency rather than memorizing the supervising DEM.

H. Dataset Characteristics

Site selection criteria We focus on mid-latitude lunar regions to ensure diverse terrain coverage. Our dataset combines multiple baseline shape models (LDEM, SLDEM, NAC DTM) with heterogeneous observations from LROC NAC and KPLO LUTI.

To systematically evaluate DEM quality across varied environments, we select eight benchmark regions widely studied across diverse lunar terrains and geological settings [3, 4, 10, 12]: the Apollo 15 (Hadley Apennine), Apollo 16 (Descartes Highlands), and Apollo 17 (Taurus Littrow) landing sites; the fresh impact craters Eimmart A and Tycho; the large sinuous rille Vallis Schröteri on the Aristarchus Plateau; and two candidate lava-tube skylights at Lacus Mortis and Marius Hills.

Regional diversity These regions span major lunar terrain types, including mare basalts, central highlands, fresh impact craters, volcanic rilles, and pit craters associated with subsurface voids. They also benefit from rich ancillary data, such as LROC stereo DTMs and LOLA altimetry, making them suitable reference sites for quantitative DEM evaluation.

H.1. Regional Characteristics and Rationale

Apollo 15 This region contains the Hadley Rille and the rugged Apennine massif, which provide strong geometric parallax due to steep local relief. The mixture of mare-highland boundaries produces sharp albedo and slope transitions, making it suitable for evaluating illumination-robust reconstruction. Extensive historical DEM products and

LOLA coverage make it a reliable reference for validating high-resolution DEM fidelity.

Apollo 16 This region represents the classic lunar highlands, characterized by highly fractured, fine-scale regolith texture and weak albedo contrast. Its low-texture surface poses challenges for both stereo-based pipelines and neural reconstruction, making it ideal for evaluating robustness in feature-poor terrain. The availability of multiple NAC orbits enables controlled assessment of geometric consistency across diverse viewing geometries.

Apollo 17 This valley exhibits complex topography enclosed by steep massifs, generating pronounced parallax suitable for volumetric reconstruction. The coexistence of dark mare soils and bright highland ejecta introduces strong photometric gradients that test shadow modeling and density estimation. Extensive Apollo-era datasets allow cross-validation against high-quality legacy topographic measurements.

Eimmart A This region provides a well-preserved crater with sharp rims and interior slopes, serving as a representative case for evaluating DEM performance in high-relief terrain. The site exhibits large orbital altitude variation across NAC passes, which naturally reduces 2D and 3D IoU and allows systematic stress testing of multi-orbit alignment, as shown in Fig. 4 of the main paper. Its fresh morphology offers clear geometric structure for analyzing reconstruction accuracy under extreme illumination changes.

Tycho The Tycho crater is one of the youngest lunar craters, with complex terracing, melt flows, and high-albedo ejecta, creating challenging photometric variation. The steep central peak and extensive ray system impose strong relief differences, beneficial for assessing depth supervision and shadow-aware rendering.

Vallis Schröteri As the largest sinuous rille on the Moon, Vallis Schröteri offers rich 3D structure with deep channels and volcanic features. Its geomorphic complexity and significant elevation drop along the rille enable testing of LNEM’s capacity to model nonlinear terrain variation. The region’s volcanic provenance complements crater-dominated sites, supporting balanced coverage of lunar surface processes.

Lacus Mortis Pit The Lacus Mortis pit features a deep collapse structure attributed to subsurface voids, providing strong cues for vertical geometry evaluation. Its sharp shadow boundaries and near-vertical walls challenge both stereo pipelines and neural rendering, particularly under

sparse-view conditions. The absence of reliable NAC DTM in pit interiors makes it ideal for assessing LNEM’s ability to fill geometric voids, as demonstrated in Fig. 6 of the main paper.

Marius Hills Pit Located within an extensive volcanic dome complex, the Marius Hills pit offers unique morphology distinct from mare and highland terrains. Although the region exhibits subtle albedo variation, the pit boundary provides strong geometric cues, making it an effective testbed for evaluating shape recovery independent of texture.

Multi-orbit observations supply both stable cross-view parallax and significantly different solar azimuth and incidence angles, which together make this site particularly suitable for assessing geometric consistency, as reported in Table 8 of the main paper, and for conducting future shape-from-shading experiments.

I. LUTI LNEM Performance and Kernel Quality Analysis

As shown in Table 5 and Table 6 of the main paper, the KPLO LUTI datasets for Apollo 15 and Tycho exhibit substantially higher $RMSE_{LOLA}$ compared to their NAC counterparts. This discrepancy primarily stems from kernel quality differences. In contrast to the LUTI experiments, NAC experiments utilize precision science ephemerides derived by the LOLA science team [8]. These smithed kernels achieve total position errors of less than 20 m, providing significantly higher geometric accuracy than the standard reconstructed kernels used for LUTI [14].

Kernel quality bounds reconstruction accuracy The Apollo 15 LUTI site achieves a bundle adjustment RMS residual of 1.38 pixels across 163 tie points with 93.8 % acceptance. This is 14 times larger than the NAC Apollo 15 residual of 0.098 pixels, as reported in Table A2.

For Apollo 15 and Tycho, SLDEM and NAC DTM both provide 1 to 2 m vertical accuracy. However, misaligned camera poses produced by reconstructed SPICE kernels cause LNEM to infer a geometrically inconsistent scene representation, resulting in large elevation errors despite accurate depth priors. This degradation is attributable to erroneous multi-view ray intersections arising from inaccurate poses, rather than limitations of the LNEM method itself. These results could be further improved with smithed kernels providing higher pointing precision [13].

Evidence for multi-view geometry learning Critically, the LUTI performance drop provides strong evidence for LNEM’s underlying learning mechanism. Despite the availability of accurate depth supervision from SLDEM or NAC

DTM, the model produces large errors because the misaligned poses of the reconstructed kernels force multi-view ray intersections to disagree. This confirms that LNEM does not merely memorize depth supervision to learn terrain geometry. Rather, the model constructs 3D structure primarily based on multi-view photometric consistency. When camera poses are inaccurate, the learned geometry diverges from the depth prior, demonstrating that our method faithfully adheres to the physical constraints of multi-view geometry. The convergence behavior in Fig. A4 and the depth supervision sensitivity results in Fig. 9 of the main paper jointly corroborate this interpretation.

J. Limitations and Future Work

Shadow modeling without physical reflectance Our shadow modeling approach relies on photometric intensity-based learning rather than the physical properties of lunar regolith. This does not fully account for secondary illumination effects and polarization characteristics that vary across terrains. Future work should incorporate physics-based reflectance models such as Hapke BRDF [2] to more accurately represent the intrinsic scattering behavior of lunar surface materials and to better generalize across sites with varying phase angles.

Mid-latitude benchmark coverage Our current benchmark focuses on mid-latitude regions. We plan to extend the dataset to include lunar south pole observations from LROC NAC, which feature extreme phase angles and persistent shadow regions that would significantly stress the shadow modeling component and provide a more complete evaluation of LNEM under challenging illumination conditions.

Multi-sensor fusion Incorporating SELENE TC imagery into Lunar Studio would enable broader multi-sensor fusion for global-scale lunar reconstruction. SELENE TC’s stereo pushbroom design and wider swath coverage complement the high resolution of LROC NAC, enabling DEM generation across diverse terrain types and latitude ranges.

Per-pixel uncertainty While σ_{total}^2 provides a natural multi-view uncertainty proxy, as shown in Fig. 8 of the main paper, explicit per-pixel uncertainty modeling near shadow boundaries would further improve the reliability of LNEM outputs for mission-critical applications such as landing site selection. Future work could explore uncertainty-aware training objectives or ensemble-based approaches to quantify reconstruction confidence more rigorously.

References

- [1] 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
- [2] Bruce Hapke. Bidirectional reflectance spectroscopy: 4. the extinction coefficient and the opposition effect. *Icarus*, 67(2):264–280, 1986. 7
- [3] Junichi Haruyama, Kazuyuki Hioki, Motomaro Shirao, Tomokatsu Morota, Harald Hiesinger, Carolyn H Van Der Bogert, Hideaki Miyamoto, Akira Iwasaki, Yasuhiro Yokota, Makiko Ohtake, et al. Possible lunar lava tube skylight observed by selene cameras. *Geophysical Research Letters*, 36(21), 2009. 5
- [4] Bradley L Jolliff, Jeffrey J Gillis, Larry A Haskin, Randy L Korotev, and Mark A Wieczorek. Major lunar crustal terranes: Surface expressions and crust-mantle origins. *Journal of Geophysical Research: Planets*, 105(E2):4197–4216, 2000. 5
- [5] Wai Chung Liu and Bo Wu. Influence of solar incidence angle on single-image photogrammetry for precision lunar topographic mapping. *ISPRS Journal of Photogrammetry and Remote Sensing*, 182:208–227, 2021. 2
- [6] 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
- [7] Roger Marí, Gabriele Facciolo, and Thibaud Ehret. Multi-date earth observation NeRF: The detail is in the shadows. In *CVPR Workshops*, 2023. 3
- [8] Erwan Mazarico, DD Rowlands, GA Neumann, DE Smith, MH Torrence, FG Lemoine, and MT Zuber. Orbit determination of the lunar reconnaissance orbiter. *Journal of Geodesy*, 86(3):193–207, 2012. 6
- [9] 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. 2
- [10] Clive R Neal and Lawrence A Taylor. Petrogenesis of mare basalts: A record of lunar volcanism. *Geochimica et Cosmochimica Acta*, 56(6):2177–2211, 1992. 5
- [11] 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. 2
- [12] RV Wagner and MS Robinson. Lunar pit morphology: Implications for exploration. *Journal of geophysical research: Planets*, 127(8):e2022JE007328, 2022. 5
- [13] Bo Wu and Wai Chung Liu. Calibration of boresight offset of Iroc nac imagery for precision lunar topographic mapping. *ISPRS Journal of Photogrammetry and Remote Sensing*, 128:372–387, 2017. 6
- [14] Youngchun Youk, Dongok Ryu, and Jeeyeon Yoon. High-resolution terrain imager development and performance evaluation for lunar exploration. *Applied Optics*, 63(1):221–227, 2023. 6

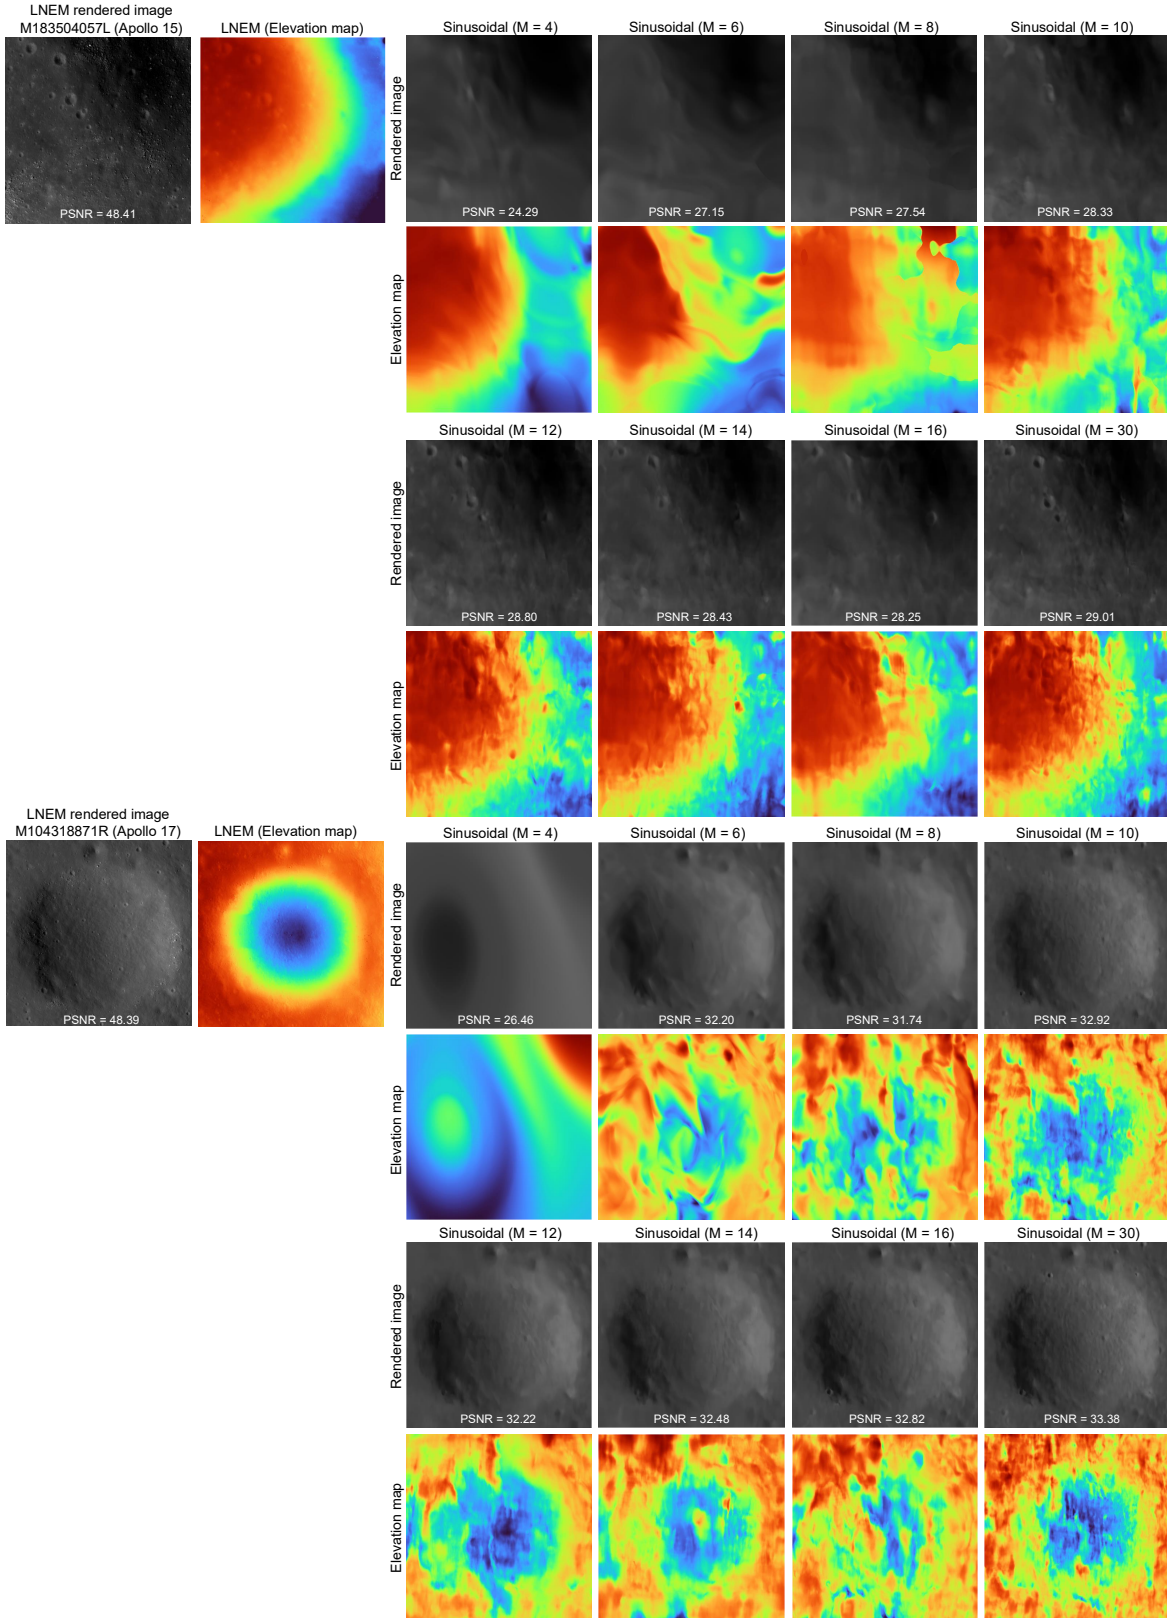


Figure A6. **Qualitative results of LNEM and Fourier feature mapping.** All results are shown at 50k iterations with identical hyperparameters, varying only the positional encoding scheme.