A Computational Framework for High-Resolution Synthetic Lunar Terrain Generation

Anders Pearson, Paolo Torrado, Joshua Smith
University of Washington
amp206@uw.edu

May 2025

Abstract

Machine learning applications for lunar exploration, such as visual navigation and radio propagation modeling, require vast amounts of annotated topographic data. However, high-resolution Digital Terrain Models (DTMs) of the lunar surface are limited. This technical report presents a computational framework for efficiently generating realistic, stochastic lunar terrains at 1 m/px resolution. Expanding upon statistical models of cratering mechanics, we implement an optimized Python-based method utilizing parallel processing to simulate geological events, topographic diffusion, and micro-roughness. We validate the synthetic terrains against Lunar Reconnaissance Orbiter Camera (LROC) data, demonstrating statistical fidelity in slope distribution and fractal roughness. The open-source software is available at: `https://github.com/anderspearson206/LunarTerrainGenerator`.

1 Introduction

The development of autonomous systems for lunar exploration relies heavily on simulation environments. Whether training reinforcement learning agents for rover navigation or estimating wireless network signal strength in complex terrain, the fidelity of the underlying environmental model is critical. While the Lunar Reconnaissance Orbiter (LRO) [1] has provided extensive data, high-resolution DTMs (1 m/px) are available for only limited regions of the lunar surface.

To address this data scarcity, we introduce a procedural generation pipeline capable of synthesizing infinite, statistically accurate lunar landscapes. Building on the cratering mechanics described by Cai et al. [2], this work introduces optimizations for computational efficiency using `Numba` [3] and extends the validation to the high-frequency spatial details required for meter-scale simulations. This report documents the algorithmic methodology and statistical validation of the generated datasets.

2 Related Works

The generation of high-fidelity lunar topography typically relies on either direct orbital observations or procedural synthesis. The Lunar Reconnaissance Orbiter Camera (LROC) has provided the most comprehensive repository of lunar topography to date, including the Wide Angle Camera (WAC) and Narrow Angle Camera (NAC) datasets [1]. However, while these instruments offer exceptional global coverage, the availability of Digital Terrain Models (DTMs) at 1 m/px resolution is very limited, which becomes an issue when large amounts of data are needed for machine learning or simulation tasks.

Procedural modeling of cratered surfaces has been explored as a means of augmenting these limited real-world datasets. Cai and Fa [2] established a robust statistical framework for modeling crater mechanics, including diameter-dependent profiles and topographic diffusion based on surface age. Although their methodology produces high-quality results at resolutions of 2 m/px to 10 m/px, the original implementation lacks the computational efficiency required to generate the large and diverse datasets needed for modern neural network training. Additionally, it is not validated against terrain at 1 m/px.

Alternative approaches have attempted to utilize Deep Learning for terrain synthesis. For example, recent research proposed a Generative Adversarial Network (GAN) for the generation of lunar DTMs at the pixel-scale by combining high-resolution monocular imagery with low-resolution DTMs [4]. While these methods can produce visually plausible textures, they often fail to maintain geometric and statistical accuracy at extremely high resolutions, frequently introducing structural artifacts. Our work bridges these gaps by optimizing the physically-grounded methods of Cai and Fa [2] for modern high-performance computing—utilizing `Numba`-based parallelization—while extending the validation to the 1 m/px scale using LROC DTMs.

3 Methods

This section details the computational algorithm for generating realistic, random lunar terrains. The methodology is grounded in established models of lunar geology and crater formation, implemented in Python with performance optimizations using libraries such as `Numba` for parallel processing. Without an atmosphere or bodies of water to erode the surface of the moon, the primary driver of topographic change is meteoroid bombardment and the resulting accumulation of impact craters. The process is divided into four stages: precomputation of crater profiles, simulation of cratering events, topographic diffusion to simulate erosion, and micro-roughness synthesis. The first three steps of this process were first presented in [2], and are adjusted in this work for efficiency. The result is an efficient and statistically accurate generator for lunar terrain for resolutions of up to 1 meter per pixel.

3.1 Crater Profile Generation

The first step in lunar terrain generation is to define the crater profile. The topographic profile, $h(r)$, where r is the radius of the crater in meters, is modeled by a piecewise function derived in [2] by examining small fresh craters from Lunar Reconnaissance Orbiter Camera (LROC) data:

$$h(r) = \begin{cases} c(D) + h_r & \text{for } \frac{r}{R} < 0.1 \\ c_3(D)\left(\frac{r}{R}\right)^3 + c_2(D)\left(\frac{r}{R}\right)^2 + c_1(D)\left(\frac{r}{R}\right) + c_0(D) + h_r & \text{for } 0.1 \le \frac{r}{R} \le 1 \\ h_r e^{-\alpha\left(\frac{r}{R}-1\right)} & \text{for } \frac{r}{R} > 1 \end{cases} \tag{1}$$

where h_r is the rim height, and α is a decay constant empirically determined to be 3.6. The coefficients $c(D), c_0(D), c_1(D), c_2(D)$, and $c_3(D)$ are diameter-dependent polynomials derived from observational data and can be found in the appendix of [2]. Given that a single square kilometer of lunar surface accumulates over 100,000 impacts, precomputing and caching crater profiles proved essential for computational efficiency.

3.2 Cratering Simulation

The crater population is initialized based on the surface age t (Ga) and the specified simulation area A_{sim} (km^2). First, the cumulative crater density for diameters $D \ge 1$ km is derived from the lunar chronology curve [5, 6]:

$$N_{\text{cum}}(D \ge 1 \text{ km}) = 5.44 \times 10^{-14}(e^{6.93t} - 1) + 8.38 \times 10^{-4}t \tag{2}$$

This reference density is extrapolated to the minimum diameter of interest ($D_{\text{min}} = 10$ m) using the Neukum production function, approximated here by a power law $N_{\text{cum}}(D) \propto D^{-b}$ with $b \approx 3.35$. The total number of craters M to be generated is then calculated as:

$$M = \left\lfloor A_{\text{sim}} \times N_{\text{cum}}(D \ge 1 \text{ km}) \times D_{\text{min}}^{-b} \right\rfloor \tag{3}$$

where D_{min} is expressed in km. Finally, the individual diameters D_i for $i = 1, \ldots, M$ are sampled via the inverse transform method [7]:

$$D_i = \left(r_i \frac{N_{\text{cum}}(D \ge D_{\text{min}})}{N_{\text{cum}}(D \ge 1 \text{ km})} \right)^{-1/b} \tag{4}$$

where r_i is a random variable uniformly distributed in $[0, 1]$. A more in-depth derivation of this procedure can be found in [2].

3.3 Topographic Diffusion

After the cratering simulation, the terrain undergoes a diffusion process to model the effects of long-term erosion and degradation. This is achieved by solving the heat diffusion equation [8]:

$$\frac{\partial h}{\partial t} = \kappa \nabla^2 h \tag{5}$$

where h is the surface elevation, t is time, and κ is the diffusivity constant. This equation is solved numerically using a finite-difference method. The diffusion process smooths sharp features, creating more realistic, aged terrain.

3.4 Micro-Roughness Synthesis and Complete Process

To further enhance realism at the target resolution of 1 m/px, the algorithm introduces high-frequency micro-roughness. This includes adding **Perlin noise**, simulating **rocks and boulders** as polygonal protrusions on the heightmap, and superimposing **small, fresh craters**. These final touches ensure that necessary high-frequency spatial details are present in the training data.

The complete process to synthesize a lunar terrain of area A_{sim} (km^2) corresponding to a surface age t is defined as follows:

1. **Pre-computation:** Precompute the crater topographic profiles for the full range of diameters of interest using Equation 1.
2. **Initialization:** Determine the total crater population and size-frequency distribution based on the specified surface age t, utilizing Equations 2-4.
3. **Time-Stepping Simulation:** Discretize the total geological timeline into intervals of Δt. For each time step:
   (a) Simulate the crater impacts occurring within this interval by sampling from the distribution and superposing their elevations onto the heightmap.
   (b) Apply topographic diffusion (erosion) to the global heightmap for the duration Δt using Equation 5
4. **Completion:** Repeat the time-stepping process until the total surface age t is reached.
5. **Post-Processing:** Apply the micro-roughness elements (noise, rocks, and small craters) to the final eroded terrain.

4 Results

While we based our work off of models derived from real data, it is important to show that these synthetic terrains accurately mimic real lunar terrain. While previous works such as [2] validated terrain generation at resolutions of 2 m/px to 10 m/px, our application requires high-fidelity data at 1 m/px. We incorporated micro-roughness elements (Perlin noise and rock distribution) and validated the output in two stages: first against a broad dataset of standard LROC Narrow Angle Camera (NAC) Digital Terrain Models (DTMs) (at resolution of 2 m/px), and second against one rare high-resolution 1 m/px DTM.

4.1 Methodology

We utilized five key topographic statistics defined in [2] to evaluate terrain similarity: Bidirectional Slope, RMS Height, Hurst Exponent, Absolute Slope, and Differential Slope.

For the broad validation, we selected 17 LROC NAC DTMs representing diverse lunar geographies at 2 m/px resolution, including Apollo landing sites, cratered highlands (e.g., Tycho), and maria regions. These large DTMs were divided into 256×256 m patches, which resulted in a dataset of over 60,000 real terrain patches for comparison against our synthetic generator.

For the high-resolution validation, we compared our synthetic output against the LROC NAC Haworth DTM, which offers native 1 m/px resolution, and is one of the lunar surface DTMs available at 1 m/px resolution. It is again divided into 256×256 m patches for comparison to our dataset.

We compare the real data to our dataset of 8,000 256×256 synthetic lunar terrains at 1 m/px resolution.

4.2 Statistical Comparison Results

4.2.1 Validation at 2 m/px Resolution

We first compared the statistical distribution of the synthetic terrains against the aggregate of the 17 LROC DTMs at 2 m/px. The results show a strong similarity across all statistics and are defined in Table 1

Table 1: Statistical Comparison at 2 m/px (Synthetic vs. 17 LROC DTMs)

Metric	Synthetic (Median [IQR])	Real (Median [IQR])
Bidirectional Slope (°)	7.40 [4.09 – 12.40]	7.90 [4.40 – 13.11]
RMS Height (m)	1.70 [0.90 – 3.09]	1.90 [1.01 – 3.16]
Hurst Exponent	0.90 [0.88 – 0.92]	0.88 [0.85 – 0.89]
Breakover Point (m)	10.0 [9.0 – 10.0]	10.0 [7.0 – 10.0]
Absolute Slope (°)	1.52 [0.83 – 2.66]	1.64 [0.90 – 2.84]
Differential Slope (°)	0.84 [0.39 – 1.61]	0.96 [0.39 – 2.55]

4.2.2 Validation at 1 m/px Resolution

To validate the efficacy of our micro-roughness synthesis, we compared the generator outputs against the 1 m/px Haworth DTM. Again, the results shown in Table 2 support the accuracy of the synthetic terrain.

Table 2: Statistical Comparison at 1 m/px (Synthetic vs. Haworth DTM)

Metric	Synthetic (Median [IQR])	Real (Median [IQR])
Bidirectional Slope (°)	7.60 [4.18 – 13.10]	7.99 [5.08 – 12.07]
RMS Height (m)	0.90 [0.49 – 1.61]	0.99 [0.65 – 1.49]
Hurst Exponent	0.898 [0.894 – 0.903]	0.898 [0.893 – 0.904]
Breakover Point (m)	13.0 [11.0 – 15.0]	12.0 [10.0 – 13.0]
Absolute Slope (°)	1.53 [0.84 – 2.67]	1.61 [1.03 – 2.46]
Differential Slope (°)	0.57 [0.26 – 1.01]	0.34 [0.12 – 0.74]

The close correspondence in topographic statistics validates that the synthetic datasets generated by this pipeline are high-fidelity approximations of the lunar environment. While there is a small difference in Differential Slope, the IQR ranges heavily overlap, and all other statistics are very similar.

5 Conclusion

In this paper, we described an efficient computational framework for generating high-resolution DTMs of lunar terrain. We demonstrated that our method produces terrains that are statistically similar to real lunar surfaces, achieving high fidelity at 1 m/px resolution across metrics including RMS height, slope distribution, and Hurst Exponent.

By leveraging `Numba`-based parallel processing and precomputed crater profiles, our framework can rapidly generate lunar DTMs. This tool provides a critical resource for the development of autonomous lunar systems, specifically for training deep learning models in visual navigation and modeling 6G wireless propagation in complex, non-terrestrial environments.

References

[1] M. S. Robinson *et al.*, “Lunar reconnaissance orbiter camera (lroc) instrument overview,” *Space Science Reviews*, vol. 150, pp. 81–124, 2010.

[2] Y. Cai and W. Fa, “Meter-scale topographic roughness of the moon: The effect of small impact craters,” *Journal of Geophysical Research: Planets*, vol. 125, no. 8, p. e2020JE006429, 2020.

[3] S. K. Lam, A. Pitrou, and S. Seibert, "Numba: a llvm-based python jit compiler," (New York, NY, USA), Association for Computing Machinery, 2015.

[4] Y. Liu, Y. Wang, K. Di, M. Peng, W. Wan, and Z. Liu, "A generative adversarial network for pixel-scale lunar dem generation from high-resolution monocular imagery and low-resolution dem," *Remote Sensing*, vol. 14, no. 21, 2022.

[5] G. Neukum, B. A. Ivanov, and W. K. Hartmann, "Cratering records in the inner solar system in relation to the lunar reference system," *Space Science Reviews*, vol. 96, no. 1-4, pp. 55–86, 2001.

[6] H. Hiesinger, C. H. van der Bogert, and N. Schmedemann, "Chronology and evolution of the lunar highlands: A new paradigm," in *Recent Advances and Current Research Issues in Lunar Stratigraphy*, vol. 482, pp. 1–52, Geological Society of America, 2011.

[7] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes: The Art of Scientific Computing.* Cambridge University Press, 3rd ed., 2007.

[8] C. I. Fassett and B. J. Thomson, "Topography of the moon's southernmost polar regions," *Journal of Geophysical Research: Planets*, vol. 119, no. 10, pp. 2255–2273, 2014.