Medium Resolution Lunar Topography: An Updated View of the Lunar Surface from Stereophotogrammetrically Derived DTMs

Jason R. Laura[1], Oleg Alexandrov[2], Lauren Adoram-Kershner[1], Kirsten Bauck[1], Benjamin H. Wheeler[1], Trent M. Hare[1]

[1] United States Geological Survey, Astrogeology Science Center, Flagstaff, Arizona U.S.A - (jlaura, bhwheeler, thare)@usgs.gov
[2] NASA Ames Research Center, Mountain View, CA, U.S.A. - oleg.alexandrov@nasa.gov

Keywords: Digital Elevatation Model, Imagery, LIDAR, Photogrammetry, Moon

Abstract

Digital Terrain Models (DTMs) are essential for lunar exploration, supporting critical applications such as landing site selection, mission planning, and in-situ resource utilization. This study presents a comprehensive reprocessing of the Kaguya Terrain Camera (TC) dataset using the NASA Ames Stereo Pipeline to generate over 130,000 medium-resolution DTMs with nearly global lunar coverage to ±70° . We observed significant bowing in the DTMs derived using the publicly available interior orientation and have used a novel technique, in the planetary sciences, to re-estimate the boresight and eight-parameter distortion coefficients. We present a jitter correction methodology to correct for vertical alignment errors in the DTMs. DTMs are aligned to the Lunar Orbiter Laser Altimeter (LOLA) point cloud and absolute errors are mean centered to zero with a mean standard deviation of 24 meters. We achieved horizontal accuracies inside the orthoimage resolution (mean of 30 meters per pixel) and 2—4 meters vertically. The methods and insights gained from this large-scale processing effort establish a framework for automating DTM derivation with other planetary data sets.

1. Introduction

Topographic data are critical enablers of science and engineering tasks in support of lunar discovery and exploration. For example, topographic products are one of the primary drivers for landing site selection where access to smoother, relatively low slope surfaces near points of scientific interest are preferable (Creech et al., 2022; Wu et al., 2022; Kirk et al., 2021; Barker et al., 2021). Likewise, surface operations, navigation, and mission planning depend upon both the topography of the local area and products derived at least in part from topographic data such as slope, surface roughness, and illumination maps (e.g., Speyerer et al. (2025); Zhong et al. (2023); Mazarico et al. (2011); Yokota et al. (2008)). Increasingly, lunar surface operations are considering the availability of in-site resources to support extended human and robotic missions. These efforts incorporate topographic data sets in the form of Digital Terrain Models (DTMS;Kumari et al. 2022). Alongside the need for resources, active research is underway simulating lunar conditions on equipment and future human operations (Kring et al., 2023). Finally, there has been an explosion in high-resolution DTM generation using shape from shading and artificial intelligence approaches (e.g., Chen et al. (2024, 2022); Liu et al. (2017); Alexandrov and Beyer (2018)). Stereoscopically derived DTMs, at lower spatial resolutions, are essential inputs and/or verification tools for newly developed and applied methods.

The Japan Aerospace EXploration Agency (JAXA) SELenological and ENgineering Explorer (SELENE) Kaguya Terrain Camera (TC) was a mapping mission designed specifically for the derivation of topography using stereophotogrammetry. The mission created and released many DTM mosaics derived from the Kaguya TC observations. There were four primary motivations to re-create a vast number of Kaguya/SELENE TC DTMs. First, the Kaguya/SELENE data remain one of, if not the best medium-resolution data sets captured using a systematic mapping orbit with near global coverage. In this work, we define medium resolution as those data larger than the 0.5 to 1.5 meter per pixel observations collected by the Lunar Reconnaissance Orbiter Narrow Angle Camera (LRONAC) and the 75 meter per pixel LRO Wide Angle Camera (WAC). These medium-resolution data are critical bridges between lower spatial resolution data sets including multi- and hyper-spectral data and high-resolution visible observations (e.g., Henriksen et al. 2017). Second, while the JAXA-released DTMs are of high quality, they are released as mosaicked tiles at a single spatial resolution with limited data provenance information (Haruyama et al., 2012). We sought to track per-DTM quality metrics, to deeply understand the processing needed to derive lunar topography such that we can provide rich quantitative and qualitative usage assessments that could be created. We also sought to create DTMs at resolutions where we could improve the size of the minimum resolvable feature visible in the released DTM mosaics. Third, in 2020 the extended mission ephemerides were updated (Goossens et al., 2020), resulting in a marked improvement in the on-ground position of the data. Finally, an initial estimate of the available spatial coverage suggested that we could increase the spatial extent of our DTMs to $\pm70°$ in most locations.

This manuscript is organized as follows. Section 1.1 describes the spacecraft, sensors, and source data used to generate lunar digital terrain models. Section 1.2 identifies the major software tools used and provides brief background on their capabilities. Section 2 steps through the process used to generate the corpus of DTMs and includes discussions of alignment of the DTM to the geodetic coordinate reference frame, re-estimation of the sensor model interior orientation, and adjustment of the exterior orientation to remove latent spacecraft jitter. In Section 3 results are presented illustrating the spatial coverage, relative accuracy when comparing overlapping DTMs, and absolute accuracy when compared to the ground truth data set. Section 4 includes a discussion of the results and pertinent lessons learned working at scale to derive lunar topography, and the manuscript concludes with Section 5 which describes potential future efforts.

1.1 Source Data

Two data sets that provide remotely sensed topographic information about the Moon are the Lunar Reconnaissance Orbiter (LRO) Lunar Orbiter Laser Altimeter (LOLA) and the SELenological and ENgineering Explorer (SELENE) Kaguya Terrain Cameras (TC1 & TC2). The Kaguya TC sensors captured over 220,000 individual observations in stereoscopic viewing mode, resulting in overlapping, fore- and aft-pointing data collection across two mission phases, a nominal mission starting September 2007 and an extended mission starting November 2008 (Kato et al., 2010).

1.1.1 Kaguya/SELENE Terrain Camera: The Kaguya/SELENE mission launched on September 14, 2007, and science operations began on October 20, 2007. The nominal or first mission phase lasted from October 20, 2007, until October 31, 2008. During the nominal mission phase, the Kaguya spacecraft flew in a circular polar orbit approximately 100 km above the surface. On November 1, 2007, an extended mission phase began. The spacecraft orbit was lowered to an average height of 50 km above the lunar surface. On June 10, 2009, the Kaguya mission ended with the impact of the spacecraft on the lunar surface. (Kato et al., 2010).

Throughout the mission, spacecraft exterior orientations were collected, post-processed, and stored in NASA Jet Propulsion Laboratory (JPL) Navigation and Ancillary Information Facility (NAIF) SPICE kernel format (Acton Jr., 1996). The SPICE or Spacecraft, Planet, Instrument, C-matrix, and Events, kernels store the interior and exterior orientation information, at some described level of quality, for a given observation from a mission. The U.S. and international planetary observing missions ubiquitously use SPICE data for spacecraft ephemerides, spacecraft-to-sensor transformations, and sensor intrinsics. For this work, we have used the JAXA-provided SPICE kernels for the nominal mission and the improved SPICE kernels created by (Goossens et al., 2020) for the extended mission. These improved kernels generally result in more accurate a priori positioning having corrected for increased gravitational perturbations, less frequent spacecraft tracking by Earth stations, and hardware issues with the spacecraft (Goossens et al., 2020). For this work, the most accurate input ephemerides were used.

The Kaguya Terrain Camera (TC) data were collected by a pair of line scan sensors looking 15° fore and aft, names TC1 and TC2. The sensors are single band operating in the 0.43 to 0.85 um wavelength and creating single band, black-and-white observations. Data for this project were collected using full (4096 pixels) and nominal (3504) pixel widths. The nominal mission phase (not to be conflated with the sensor nominal observing mode) collected data that were typically 35 km wide and 100 km long, at an approximately 10 meter per pixel ground sample distance. The mission team clipped these long, slender swaths to a uniform length for release. The extended mission phase, with a lower orbit, results in observations closer to 30 km in total width with a ground sample distance of 5—8 meters per pixel. Throughout the mission, the Kaguya TC sensors imaged over 99% of the lunar surface. The Kaguya TC sensor operated in three modes during the mission. Of interest to this work is the stereoscopic observing mode where the fore- and aft-pointing sensors collected data concurrently with 30° of parallax, Figure 1. Throughout the nominal and extended missions over 220,000 stereoscopic observations were acquired (110,000 stereo pairs). These data have near-global coverage in our region of interest of ±70° where extreme solar illumination angles are less of an issue.

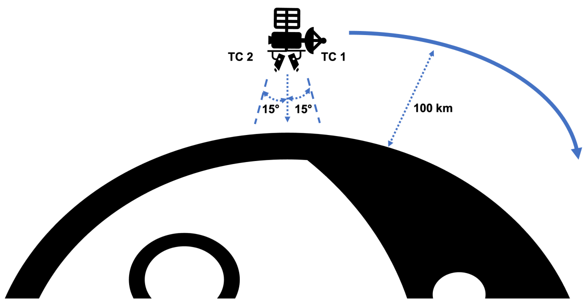

Figure 1. The Kaguya TC sensors, TC1 and TC2, during the nominal mission phase.

1.1.2 Lunar Reconnaissance Orbiter Lunar Orbiter Laser Altimeter: While it is possible to create an internally consistent, relatively controlled (Laura et al., 2018) data set using only the Kaguya TC data, these data would effectively float, with some error, over the currently defined geodetic coordinate reference frame. To absolutely constrain the Kaguya TC data to the adopted reference frame, Moon2015 Mean Earth/Polar Axis, as defined by the International Astronomical Union (IAU; Archinal et al. 2018), one must tie to some reference data set. The Lunar Reconnaissance Orbiter (Tooley et al., 2010) Lunar Orbiter Laser Altimeter (LOLA;Smith et al. 2011) instrument, which has collected over 6.3 billion spot altimetry readings since July 2009 provides such a data set. The LOLA instrument is a multibeam laser altimeter that illuminated 5 spots on the ground with each spot having a 5-meter diameter (assuming nadir pointing, no surface slope, and a nominal 50-kilometer spacecraft altitude) (Smith et al., 2011).

1.2 Software Tools

We utilized three software libraries on the United States Geological Survey’s (USGS) Hovenweep supercomputer (Falgout et al., n.d.) to create digital terrain models (DTMs). These libraries provided data pre- and post-processing capabilities, stereo photogrammetry capabilities, and data interoperability and discovery tools.

1.2.1 Integerated Software for Imagers and Spectrometers: The Integrated Software for Imagers and Spectrometers (ISIS) software library is maintained by the USGS Astrogeology Science Center. ISIS is the premier library for planetary data ingestion, photogrammetric control, map projection, and interactive data analysis. We have used the ISIS ingestion tool *kaguyatc2isis* to read the archived observations into an interoperable format and the Abstraction Library for Ephemerides (ALE) (Mapel et al., 2019) library to generate the necessary interior and exterior orientation for each observation. ISIS also provides the Community Sensor Model (CSM; described below in Section 1.2.3) for the Kaguya TC sensor.

1.2.2 NASA Ames Stereo Pipeline: The National Aeronautics and Space Administration (NASA) Ames Stereo Pipeline (ASP) provides four broad categories of algorithms, via command line tools, that support the creation of stereo photogrammetrically derived DTMs (Beyer et al., 2018). We

https://doi.org/10.5194/isprs-archives-XLVIII-M-6-2025-197-2025 | © Author(s) 2025. CC BY 4.0 License.

classify these as computer vision techniques, stereophotogrammetry capabilities, point cloud manipulation and alignment algorithms, and geospatial processing tools. The tools that ASP provides for several different command line tools that can be used to apply standard computer vision algorithms to images and stereo pairs. The computer vision tools support interest point detectors, image interest point correlators to find image correspondences, and outlier detection algorithms to remove spurious correlations. Once correlations are found, ASP can build a photogrammetric control network (for two or more images) and perform a standard block bundle adjustment (using the CERES solver) to achieve n-image alignment. ASP also provides tooling to perform dense image-to-image matching (for each pixel in the arbitrarily selected 'left' image), epipolar alignment, subpixel alignment, outlier rejection, triangulation, and estimation of jitter for removal. These algorithms are the stereophotogrammetric methods needed to take two observations, with associated sensor models, and yield a three-dimensional point cloud. Next, ASP offers an implementation of the Iterative Closest Point (ICP) algorithm (Pomerleau et al., 2013) via the *pc_align* tool to apply a linear transformation (a 3D affine transformation) to align two point clouds. Finally, necessary geospatial transformations are provided to grid point clouds, map project results, and perform data format conversions. ASP provides the majority of the processing tools used in the completion of this work.

1.2.3 Community Sensor Model: The Community Sensor Model (CSM; Group 2007) is a widely adopted specification and associated application programming interface (API) that standardizes photogrammetric operations across heterogeneous sensor types and ephemerides sources. CSM-compliant sensor models are used throughout this work for all photogrammetric operations because the CSM provides programmatic access to low-level photogrammetric functions like image-to-ground and ground-to-image, and high interoperability across the software libraries we used.

A functioning CSM sensor model is composed of two things (Laura et al., 2020). First, a generic implementation, for example, a line scan sensor, or framing sensor that implements the CSM specification. Second, the interior and exterior orientation that are associated with a given observation. The interior and exterior orientation is encoded in either an Image Support Data (ISD) file or a sensor state file; both contain the same information, just in different machine-readable formats. Of particular interest in this work are the sensor intrinsics, namely: focal length, principal point offset, and the four distortion coefficients that describe the Kaguya TC sensor model. Each of these intrinsics is defined as an adjustable parameter allowing one to use a bundle adjustment to estimate corrections.

2. Methods

Herein, we describe the process of generating Kaguya TC DTMs with nearly global (to ±70°) coverage. The creation of the DTMs was an iterative process whereby an initial set was generated and assessed for quality. This assessment discovered systematic issues with the sensor model and we describe the novel approach used to solve for and update the Kaguya TC sensor model intrinsics across both sensors, observing modes, and observing time of day. This section is organized linearly, first describing the initial DTM derivation, then the issues discovered, the methods employed to correct the issues, and finally, the full pipeline as used to generate the released DTMs.

2.1 Initial Stereophotogrammetric DTM Derivation

As described above, the ∼100-km-long-swaths are clipped into scenes for delivery to users. Each scene, within a swath, is 4656 pixels in the along-track (north to south or south to north) length, or 37—45 km on the ground. Metadata accompanying each observation includes an orbit number, scene number, and view geometry information. We first winnow the ∼220,000 and remove observations that are 70° or more off-nadir. Next, stereo pairs need to be identified. To accomplish this, we find pairs with the same orbit number (revolution), strip number (identifying which ∼100 km strip the data are in), and adjacent scene number. Having these candidates, we finally check that the stereo pair contains one observation each from TC1 and TC2 (the fore- and aft-facing sensors). Outside the winnowing for oblique observations, we do not concern ourselves with the view geometry as the stereoscopic observing mode ensures that observations are captured with a nominal 30°parallax angle. Further, the identification of stereo pairs by orbit number and scene number ensures that data are collected with, at worst, very similar solar illumination.

Each stereo pair is independently preprocessed. First, the JAXA archived data are converted into the CUB format using the ISIS *kaguyatc2isis* command and the appropriate SPICE data or ephemerides are attached using *spiceinit*. The *spiceinit* application writes ephemerides as binary blobs alongside the observation data. To create a CSM sensor model, the ALE *isd_generate* command line tool, that ships with ISIS, is used to extract the observation interior and exterior orientation and write a CSM-compliant ISD.

Next, ASP is used to perform a two-image block bundle adjustment. The block bundle adjustments first use the SIFT algorithm to identify the initial image correspondences and RANSAC to discard outliers. The adjustment is performed to bring the images into good relative alignment before matching. We find that performing an initial relative adjustment produces higher quality DTMs with less noise and null value pixels than using only the a priori exterior orientation. The result of this step is a series of adjustments to the extrinsics of the left and right images.

A low-resolution DTM is then generated to support the orthorectification of the left and right images before high-resolution DTM generation. This step is needed to improve the overall quality of the image matching and is a key step in generating a high-quality DTM. The ASP *parallel_stereo* is then used to generate a low-resolution DTM. First, the left and right images are tiled for parallel processing, and image normalization is run to bring the observations into close radiometric parity. Then, the SIFT algorithm is used to detect 400 interest points per tile, a normalized cross-correlation algorithm is run to achieve pixel scale alignment, and RANSAC is applied to remove outliers. Next, an affine epipolar transformation is performed to bring the left and right images into rough alignment (Beyer et al., 2018). The result of this step is an initial disparity map that encodes the integer pixel-wise differences between the left and right images. Subpixel alignment is then achieved using a Bayes expectation-maximization (EM) weighted affine-adaptive window correlator (Broxton et al., 2009; Nefian et al., 2009; Beyer et al., 2018). As described by Beyer et al. (2018), the adaptive search template used by the Bayes EM provides significantly improved subpixel correlation in conditions with high slope or disparate view geometries. Finally, the point cloud is triangulated by computing the closest point intersection of the look vectors of each shared point in the left and right images and a gridded, low-resolution DTM is generated. The *point2dem* tool is then used to grid the derived point cloud to a 400-meter-per-pixel-product, with holes smaller than 50 meters filled using bicubic interpolation, and map-projected using a DTM-centered orthographic projection.

A high-resolution DTM is generated by first orthorectifying the left and right images using the previously created low-resolution DTM. The process of image correspondence identification, outlier detection, subpixel refinement, and triangulation is then repeated with a few small alterations. First, an epipolar transformation is not applied to the left and right images because they are already map-projected and orthorectified. Next, in addition to the SIFT computed image correspondences, we also generate a set of gridded image correspondences for later use when spacecraft jitter is solved from the exterior orientation. Finally, we do not grid the point cloud to a final DTM product. The high-resolution point cloud is only relatively accurate to the lunar geoid and an alignment must be performed for absolute alignment and subsequent accuracy assessments.

As stated above, one goal during the generation of these DTMs was to improve upon the effective resolution of the DTMs by decreasing the size of the minimum resolvable feature (MRF) while controlling for the number of spurious artifacts introduced. To that end, a round of parameter estimation tests were run varying the image correlation algorithm, the kernel sizes, the alignment methods, and the final DTM gridding resolution. Ultimately, block matching with the Bayes EM subpixel correlator algorithm and an affine epipolar alignment in the low-resolution DTM generation processing step were selected as providing the most consistent results. The key parameter for Kaguya TC, particularly those data with lower illumination angles is the use of Bayes EM subpixel alignment.

The LOLA point cloud data define the Moon2015 coordinate reference frame (Archinal et al., 2018) and are assumed to be ground truth. The LOLA data are freely available in the Cloud Optimized Point Cloud (COPC) format and easily queryable via an API to subset for spot observations in a given region of interest. Therefore, we query LOLA in the approximate region of each DTM, buffering the footprint by 0.1 degrees in latitude and longitude to account for inaccuracies in the exterior orientation. The ASP *pc_align* tool is then used iteratively to compute the transformation from the sparse LOLA point cloud, Figure 2 for an example of the data, to the dense DTM that has been generated. *pc_align* using an iterative closest point algorithm to compute the 3D affine transformation between two point clouds. We run this algorithm up to three times on each DTM, first with an initial guess of the threshold beyond which outliers should be removed. If that threshold is too small, as reported by the tool, we re-run the initial gross alignment parametrizing using the tool’s reported error threshold. After gross alignment, we then use the non-ICP *Fast Global Registration* algorithm to apply small adjustments. The results are an affine transformation, and an inverse transformation, usable to align the DTM with LOLA.

The aligned point cloud is then post-processed to extract a gridded DTM, the exterior orientation stored in the left and right image CSM state files are updated using the DTM to LOLA 3D affine transformation, and accuracy metrics are computed. To grid the DTMs, we make use of the standard convention whereby final DTM post spacing is computed by multiplying the largest input ground sample distance of the stereo pair by

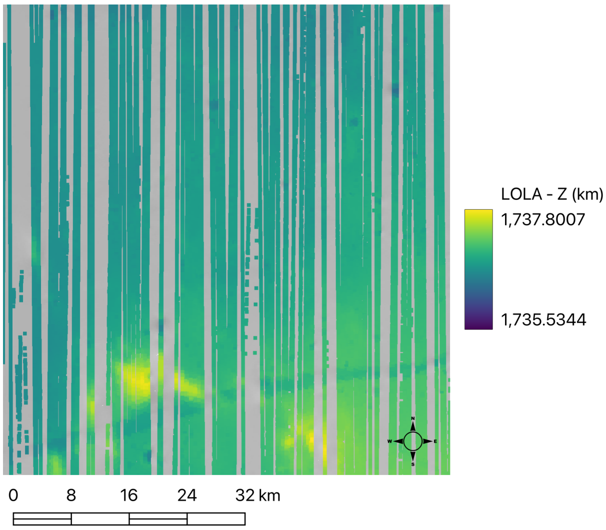

Figure 2. LOLA tracks in a ∼60 kilometer wide region of interest at the equator. Gaps as large as 4 kilometers can be measured. The LOLA point cloud provides the ground truth to which all derived DTMs are aligned.

three (Kirk et al., 2021). While this general rule is not perfect, it does provide a widely adopted guideline that we have opted to use.

2.1.1 Initial Results: A large number of DTMs, spatially distributed across the lunar surface, were then generated. For each stereo pair a DTM, left-orthoimage, stereo pair intersection error raster, and summary error statistics text file were available for quality analysis. Significant bowing in overlapping DTMs was observed. In Figure 3, the offsets between the LOLA point cloud and the stereo photogrammetrically derived point cloud are shown. Significant and systematic ”U” shapes bowing is present in every DTM.

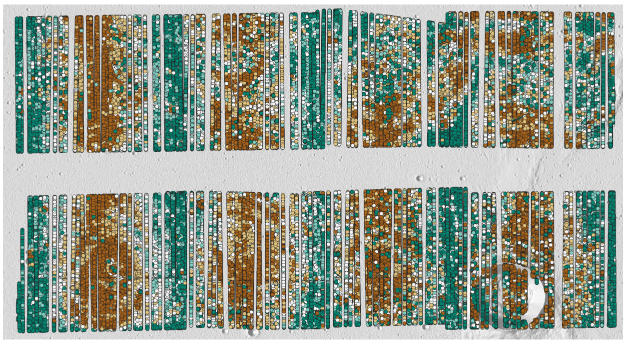

Figure 3. Each point represents a delta between a stereoscopically derived DTM and the ground truth LOLA point cloud. These two rows, each with four DTMs are visible with the center of each DTM generally being defined by a dark brown stripe of points, where the topography is underestimated by approximately 20 meters and the edges are defined by dark green sets of points where the topography is overestimated by approximately 20 meters. Each DTM exhibits a pronounced ”U” shaped bowing.

2.1.2 Solving for Updated Intrinsics: Using the results above, we hypothesized that errors in the interior orientation, particularly that the eight parameters (4 x-coefficients and 4 y-coefficients) tangential distortion model and the boresight were causing the observed bowing effect. Good horizontal alignment was observed between overlapping DTMs strengthening the assessment that systematic errors in the exterior orientation were not the cause. Additionally, systematic errors in the intersection error raster, generated during triangulation, were observed. Finally, after re-estimating the interior orientation for a single observing mode and time of day we noted that the observed errors were corrected in only a subset of the data. We determined that the interior orientation needed to be re-estimated for data with different observation modes (e.g., nominal or full swath width) and data acquisition times of day (e.g., morning or evening).

A five-step process was applied to re-estimate the interior orientation. First, 10 stereo pairs were selected with the same swath width, the same illumination or acquisition time, and the left image being acquired by TC1 and the right image by TC2. Each stereo pair had approximately 30% or more overlapping areas with adjacent stereo pairs in our test set. Next, dense interest points were detected along a regular grid, and a block bundle adjustment was performed on all twenty of the input images with the interior orientation fixed. This brought all of the observations into relative alignment with one another by updating their exterior orientations. Third, the computed reprojection errors were observed across all stereo pairs to ensure that the magnitudes were low and that no systematic spatial effects along image edges were observable. Topography-induced reprojection errors were still evident, but these did not impact our final results as they are random concerning the observation coverage. Fourth, ten independent DTMs were created using the previously described stereo triangulation and gridding techniques. These DTMs were created at the 10-meter per-pixel image resolution. The DTMs were manually inspected for inter- and intra-DTM quality and compared to the LOLA point cloud. The highest achievable accuracy was important in this step. Finally, a bundle adjustment was performed using the previously computed dense interest points, the previously updated sensor extrinsics that brought the 20 images into relative alignment. The height of each point, using the adjustment, was extracted from the high-fidelity DTM created in the previous step. Error in the extrinsics is attenuated by using the updated exterior orientation and heights from the high-fidelity DTM computed in the previous step. Effectively, the exterior orientation is highly accurate and the exterior orientation is allowed to vary.

The interior orientation has been re-estimated for the TC1 and TC2 sensors for the following observing modes and conditions: (1) full swatch width, (2) nominal swatch width morning, and (3) nominal swath width evening. Testing demonstrated that the same interior orientation values were for usable for both morning and evening full swath width observations.

2.1.3 Latent Spacecraft Jitter:

Spacecraft jitter is caused by high-frequency errors in the exterior orientation and manifests as systematic distortions or ripples, perpendicular to the data acquisition path. These errors are most prominent in triangulation error raster or derived hillshade. Figure 4 shows the intersection error before and after updating the sensor's interior orientation. One can see the systematic red (high error) and blue (low error) striping perpendicular to the data acquisition path. The elliptical shape in the before images is a combination of error in the interior orientation and latent spacecraft perturbation. After updating the interior orientation, one sees that the error is now rectangular in shape and attributable primarily to spacecraft jitter.

The final processing step involved solving for spacecraft jitter

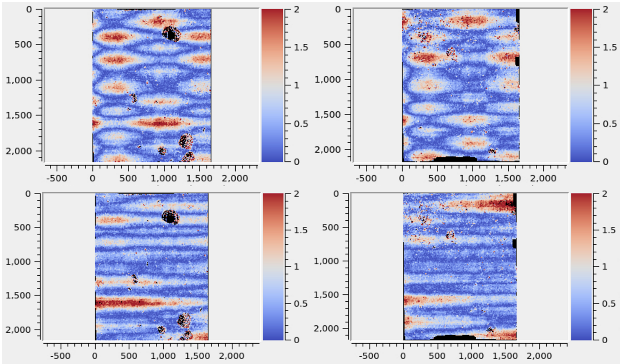

Figure 4. Observed triangulation error for two DTMs (columns) where the first row uses uncorrected interior orientation and has a clearly visible elliptical error pattern (red) and the second row uses the corrected interior orientation and shows latent, uncorrected spacecraft jitter (reds to blues). The black areas have no data values. Color bar units are meters.

by adjusting each position and orientation in the exterior orientation to a known ground, (e.g., LOLA). First, the LOLA point cloud used for alignment was gridded to a 25-meter-per-pixel product. This is a significant oversampling from what one would use for a science product but is appropriate as a ground constraint for this work. Next, the ASP *jitter_solve* was run to solve for spacecraft jitter. Jitter correction is formulated as an optimization problem. For this work, the interior orientation is constrained such that little variability is allowed. This is because spacecraft intrinsics have already been solved and the images exhibit good alignment. The tool constrains the intersected points from each input image to the provided lower-resolution DTM, in this case, LOLA. We know that the generated DTMs are in good horizontal and vertical alignment because of the previous processing steps. Therefore, we set a 10-meter uncertainty on the intersected heights. Figure 5 illustrates the triangulation error before (left) and after (right) correction. We note that not all data can be successfully jitter-corrected, but most can.

2.1.4 Pipeline with Corrected Intrinsics:

The full data set was then processed using the above pipeline. We note that as of the time of writing, we are re-processing DTMs with the added spacecraft jitter correction step and expect to have those released in early 2025. To recap the process. We first preprocess the data using ISIS and modify the interior orientation to use the distortion coefficients and boresight that we computed for this work. We then generate both low- and high-resolution point clouds using ASP. The high-resolution point cloud is then aligned to LOLA using the ICP and FGR algorithms. A spacecraft jittering correction step is then run to attempt to correct for any latent spacecraft movement during image acquisition. Finally, a gridded DTM is created, accuracy metrics are computed, and ancillary data products are staged for release as analysis-ready data.

3. Results

In total, 130,799 Kaguya TC DTMs were created and released using the above-described pipeline. At the time of writing, only a handful have undergone the final, jitter correction step; processing is ongoing pending the availability of computing resources. The released DTMs provide almost complete coverage in the study ±70° study area. Figure 6 illustrates both the

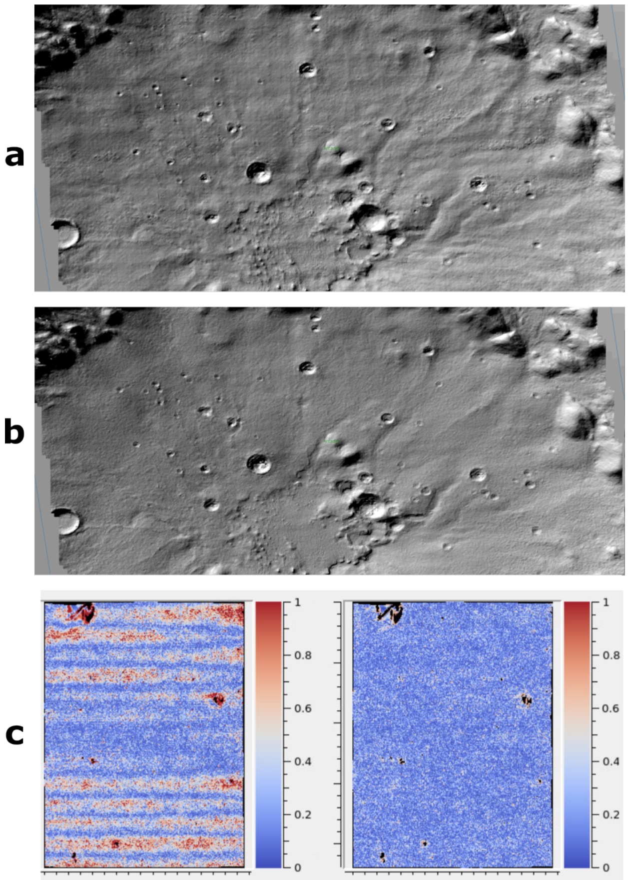

Figure 5. (a) Hillshade generated without spacecraft jitter correction. High-frequency, across-track errors are visible. (b) Hillsahde generated after jitter correction showing significantly reduced high frequency, across track artifacts. (c) Before (left) and after (right) triangulation errors show a near-total correction of the latent spacecraft jitter. Color bar units are meters.

coverage and the final achieved spatial resolutions. The spatial autocorrelation in the resolutions, visible as curved bands across the surface, and the clustering of lower resolution DTMs poleward are expected as they mirror the orbital characteristics of the spacecraft. We also note that the region centered around -170° longitude, -55° latitude is particularly sparse with much of the nominal mission data in that region failing to process.

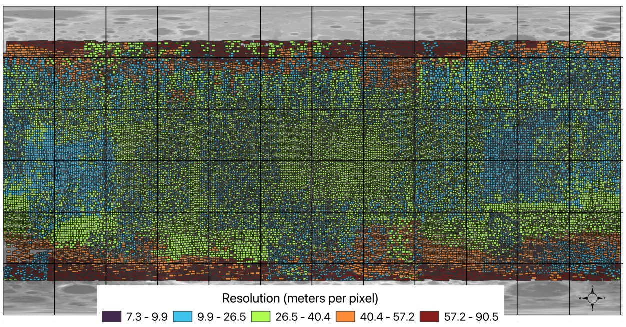

Figure 6. A total of 130,799 DTMs spanning -180° to 180° , shaded based on final product resolution. A 30° by 30° graticule overlays the data. The base map is the SLDEM2015 product created by Barker et al. 2021 using the JAXA-derived Kaguya TC DTMs.

A goal of this work was to achieve the smallest possible minimum resolvable feature, using stereophotogrammetric techniques. In Figure 7 we report the distribution of ground sample distances for the final DTMs. The highest resolution DTM released has a GSD of just over 7 meters per pixel. The mean resolution for all DTMs is 32 meters per pixel and the highest resolution DTM is 90 meters per pixel. Visual inspection of the input images and derived hillshades shows that features on the order of 3—5 pixels are resolvable. Therefore, across all DTMs, we can resolve features between 21 and 270 meters in diameter (measuring primarily craters). We believe that further improvements will require using different processing techniques such as shape-from-shading (e.g., Alexandrov and Beyer (2018)). Given that the bulk of the input were collected during the nominal mission with an average 10 meter per pixel GSD, the final 30 meter per pixel DTM GSD is as expected.

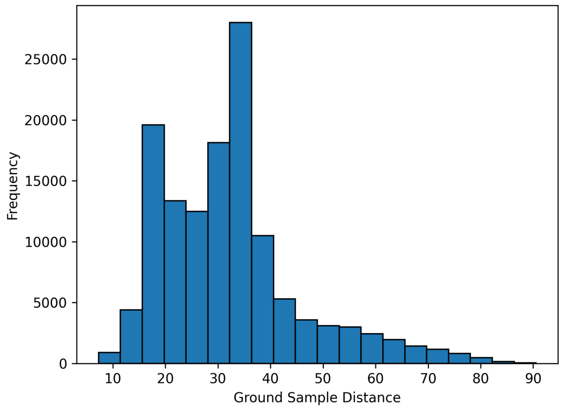

Figure 7. A histogram showing the ground sample distribution of the generated DTMs using $\sim$4 meter bin widths.

Figure 8 shows a merge of over 120 DTMs over an arbitrarily selected region of interest. All DTMs share the same color ramp. Inter-DTM vertical alignment varies spatially. Measurements away from the absolute edges of any pair or trio of DTMs show sub-meter vertical errors. Measurements at the edges of one DTM show errors on the order of 4 to 6 meters vertically which is a significant improvement over the previously measured 20 meter vertical errors. We attribute these edge effects to latent errors in the interior orientation and residual spacecraft jitter that have not been accounted for. A comparison of the hillshade and orthoimages horizontally is described below.

Figure 9 shows the final offsets between LOLA, the ground truth, and the Kaguya TC-derived DTMs. We see that the systematic, across-track errors present in Figure 3 have been removed and that the error is generally randomly distributed, save some latent spatially autocorrelated errors that align with lunar topography (e.g., in the northeastern pair of DTMs). This is due to under or overestimation of the topography in the stereo matcher and is attributable to more challenging illumination conditions. Globally, *pc_align* has worked as expected and the mean, per-DTM vertical offset to LOLA is mean centered to effectively zero (0.11 observed, n=130.799, no outlier DTMs removed). More informative, at the 99.5% interval, we report a mean standard deviation of 4.3 meters to LOLA, indicating that these DTMs have achieved global vertical alignment inside the anticipated error budget.

Errors in the sensor extrinsics translate directly to errors in the

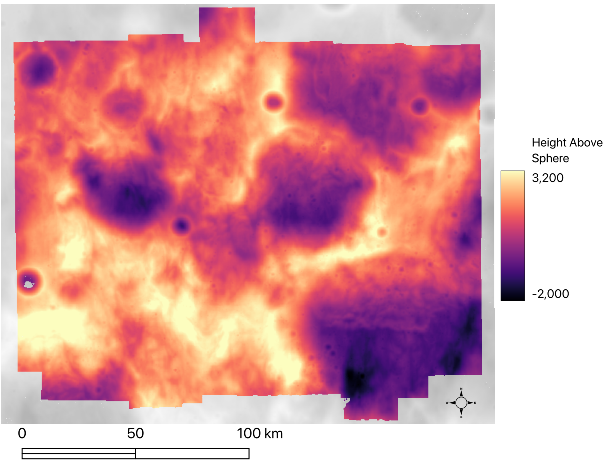

Figure 8. Over 120 DTMs with matched color ramps illustrating visual vertical consistency.

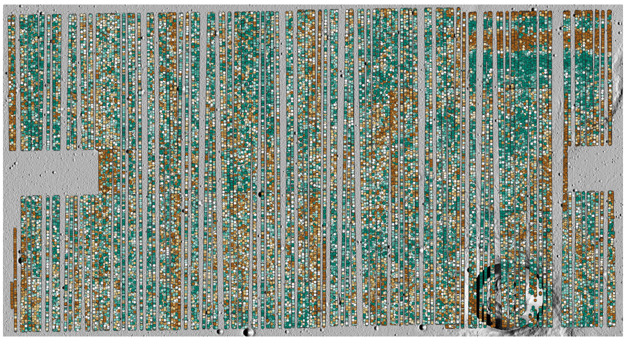

Figure 9. Observed per-point offsets between the LOLA ground truth point cloud and the final, non-jitter corrected DTMs. Dark browns are positive 4-meter errors and dark greens are negative 4-meter errors (class breaks are standard deviations from the zero mean.)

absolute positioning of the final DTM products. Figure 10 presents the North, East, and Down adjustments to each DTM applied by *pc_align* to minimize the error between our derived point cloud and the LOLA point cloud. We see the majority of the adjustment occurring along-track, in the north direction with a median adjustment of just over 250 meters and a mean adjustment of 138 meters. The across-track shifts are significantly lower with a mean adjustment of -38 meters and a median adjustment of zero. Finally, in the z-direction, the mean adjustment is 16 meters and the median is 20 meters. The mean and median magnitude of adjustment, aligning the derived point clouds with LOLA is 325 meters. When coupled with the observed adjustment standard deviations in each dimension, 253, 187, and 21 meters respectively, the DTMs appear, at least numerically to be consistently adjusting to LOLA.

To test our assessment that the point cloud adjustments to LOLA, while independently applied, resulted in a consistent adjustment, an orthomosaic was generated. Figure 11 shows the result of this test in one region of interest. Here 67 individually aligned orthoimages were mosaicked, taking the average pixel value for each grid cell. Using this averaging approach, one expects to see blur or softness where two or more images are out of alignment. Here, we see good inter-orthoimage alignment which we extrapolate to be indicative of good horizontal DTM alignment. Coupled with the vertical accuracy results presented above, we are confident that these DTMs are both internally consistent and in the best possible horizontal and vertical alignment to LOLA.

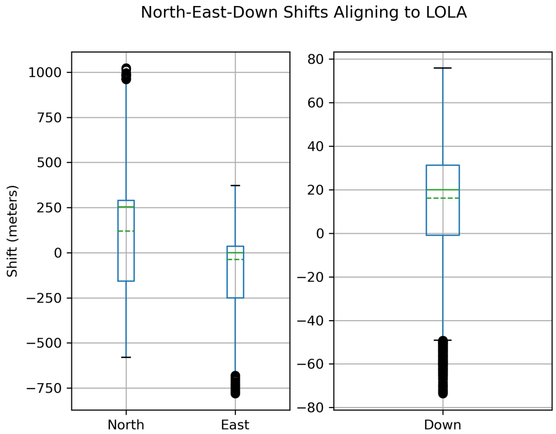

Figure 10. These boxplots show the North-East-Down shift magnitudes for 99.5% of the generated DTMs. The remaining DTMs are assessed to have erroneous shifts and are discarded. Median lines are solid green and mean lines are dashed green.

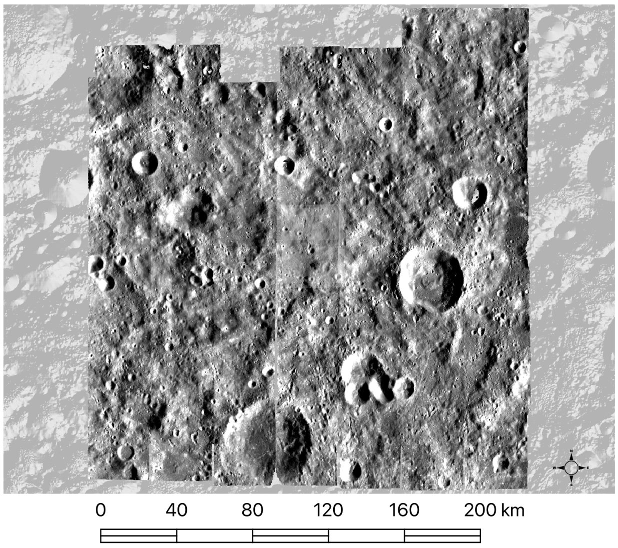

Figure 11. A mosaic was created using the average of each pixel from 67 overlapping orthoimages. We see subpixel registration between the orthoimages. Mosaic is shown over a hillshade derived from the 60-meter per pixel SLDEM2015 product Barker et al. (2016).

4. Discussion

The results presented demonstrate that we have achieved the goals that we set out to accomplish. First, we have created and publicly released over 130,000 individual stereo photogrammetrically derived DTMs using the Kaguya TC stereoscopic data. Each DTM released includes a provenance file that tracks all processing, from ingestion into ISIS through processing using ASP. These provenance files ensure the reproducibility of this work and allow users to use our DTMs with confidence. Second, the corpus of DTMs released includes a sizable number generated from extended mission data. Those data, in general, have a smaller GSD because of the lower spacecraft orbit, Figure 7, resulting in smaller post-spacing in the final derived DTMs. In addition to the goals we had when initiating this work, we also identified and corrected errors in the intrinsics, errors caused by spacecraft jitter, and latent errors in alignment to LOLA caused by previously uncorrected errors in the extrinsics.

The vertical and horizontal accuracy of these data makes them a key bridge product, usable to connect the highest resolution lunar data to LOLA, the coordinate reference frame proxy. In instances where high-resolution data are directly tied to LOLA, these Kaguya TC DTMs are also quite valuable, as they are easily mosaickable and usable for investigations requiring larger spatial extents (e.g., large feature AI/ML training, geologic mapping, large-scale change detection).

In addition, we find that automated DTM generation at this scale is possible assuming several considerations are met. First, systematic processing of a subset of the data is a necessity before any large-scale processing can occur. One must ensure that the sample data from multiple spatially distributed and representative regions of interest are selected. These data must span the mission and account for any known temporal differences in data quality (e.g. when a star tracker stopped functioning and position inaccuracies are increased). These data must also span the observing modes and conditions under which data were collected. Without testing morning- and evening-observing times as well as nominal- and full-swath width data, we may not have identified the systematic errors in the intrinsics before large-scale processing.

While this seems intuitive after performing this work, it is important to remind oneself that good horizontal alignment between overlapping observations is not adequate to rule out systematic errors in the interior orientation. In this work, we observed both good inter-image alignment at resolution and systematic 20-meter vertical errors across DTMs. Therefore, we suggest that, when possible, planetary sensor model developers consider the generation of DTMs as an essential validation and verification step.

Next, the ICP algorithm implemented by the ASP *pc_align* tool is essential in aligning DTMs to LOLA, the accepted ground truth. One must use caution with this tool and diligently inspect the ancillary reporting files created by ASP. We see that *pc_align* does a terrific job mean centering the overall DTM error around zero. The descriptive statistics provided by ASP, via *geodiff* are then one quickly accessible indication of the overall quality. A visual inspection of the per-point differences between the derived DTM and LOLA is essential to ensure that a mean zero error DTM is not suffering from systematic across-track errors, or long- or short-wavelength along-track errors.

We were surprised to find that at times, latent spacecraft jitter can be hard to distinguish from errors in the sensor model intrinsics. For example, we attempted to solve for the interior orientation for the full swatch width morning data and observed no improvement in the vertical accuracy to LOLA. For this reason, the same interior orientation is used for all full swath width data, Section 2.1.2. For those data, only the correction for latent spacecraft removed the systematic observed errors.

Finally, this work demonstrates the need to develop a robust set of quality assurance criteria that can be universally applied to planetary DTMs where in-site accuracy assessments are largely impossible. These techniques need to support the assessment of the data for local inaccuracies and blunders, long and short wavelength errors due to spacecraft jitter and other effects, and qualitative assessments that capture the relative connectedness or topology of valid data.

5. Future Work

There are several promising avenues for future work. First, we are currently working to fill any remaining holes, data permitting, in the ±70° region of interest. To achieve adequate coverage we will have to be less stringent in our DTM selection criteria and vary parameterization for data with starker illuminations. Next, as seen in Figure 11, it is possible to take the DTM orthoimages and create regional orthomosaics. It is also possible to expand that process to create a globally contiguous orthomosaic. We believe that a mean of 30 meters per pixel is the most appropriate scale, balancing some supersampling at the higher latitudes with downsampling at the mid-latitudes. This product would be terrific for geologic mapping at certain scales and as a base map for context and interactive visualizations. We note that all of the orthoimages are currently available and that regional mosaics can be rapidly created today using one's GIS of choice. Third, we are working to take the updated sensor ephemerides and create an at-resolution image mosaic. Each DTM is aligned to LOLA and we demonstrate good inter-DTM alignment in this work. As part of the alignment, we have created updated exterior orientation files for each left and right observation in the stereo pair. Therefore, we have effectively controlled each input image to LOLA. We estimate that the accuracy of that alignment is inside of an orthoimage pixel, or approximately 30 meters. Initial tests have demonstrated that it is possible to use the updated ephemerides, a network of image correspondences, and a standard block bundle adjustment to further adjust the observations to bring them into subpixel alignment. This adjustment can be constrained to maintain the absolute alignment to LOLA and results in at-resolution inter-image alignment. More work is required here to build the many image control networks needed to globally apply this process. Fourth, we will be releasing a merged body-centered body-fixed point cloud merging all of the triangulated points from each DTM. We anticipate that this product will allow users to make their determination about appropriate grid spacing if and when they need to convert from a point cloud to a gridded product.

6. Acknowledgements

We would like to thank the Kaguya/SELENE team at JAXA for collecting these data and making them publicly available, Dr. Hiroyuki Sato for the valuable discussion around the data, and Dr. Laszlo Kestay for the creation of Figure 1. This work is supported by NASA grant #NNH22OB02A. Any use of trade, firm, or product names is for descriptive purposes only and does not imply endorsement by the U.S. Government.

References

Acton Jr., C. H., 1996. Ancillary data services of NASA's Navigation and Ancillary Information Facility. *Planetary and Space Science*, 44(1), 65 – 70.

Alexandrov, O., Beyer, R. A., 2018. Multiview Shape-From-Shading for Planetary Images. *Earth* *and Space Science*, 5(10), 652–666. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1029/2018EA000390.

Archinal, B. A., Acton, C. H., A’Hearn, M. F., Conrad, A., Consolmagno, G. J., Duxbury, T., Hestroffer, D., Hilton, J. L., Kirk, R. L., Klioner, S. A., McCarthy, D., Meech, K., Oberst, J., Ping, J., Seidelmann, P. K., Tholen, D. J., Thomas, P. C., Williams, I. P., 2018. Report of the IAU Working Group on Cartographic Coordinates and Rotational Elements: 2015. *Celestial Mechanics and Dynamical Astronomy*, 130(3), 22.

Barker, M. K., Mazarico, E., Neumann, G. A., Smith, D. E., Zuber, M. T., Head, J. W., 2021. Improved LOLA elevation maps for south pole landing sites: Error estimates and their impact on illumination conditions. *Planetary and Space Science*, 203, 105119.

Barker, M., Mazarico, E., Neumann, G., Zuber, M., Haruyama, J., Smith, D., 2016. A new lunar digital elevation model from the Lunar Orbiter Laser Altimeter and SELENE Terrain Camera. *Icarus*, 273, 346 – 355.

Beyer, R. A., Alexandrov, O., McMichael, S., 2018. The Ames Stereo Pipeline: NASA’s Open Source Software for Deriving and Processing Terrain Data. *Earth and Space Science*, 5(9), 537–548.

Broxton, M. J., Nefian, A. V., Moratto, Z., Kim, T., Lundy, M., Segal, A. V., 2009. 3D Lunar Terrain Reconstruction from Apollo Images. G. Bebis, R. Boyle, B. Parvin, D. Koracin, Y. Kuno, J. Wang, J.-X. Wang, J. Wang, R. Pajarola, P. Lindstrom, A. Hinkenjann, M. L. Encarnação, C. T. Silva, D. Coming (eds), *Advances in Visual Computing*, Springer Berlin Heidelberg, Berlin, Heidelberg, 710–719.

Chen, H., Gläser, P., Hu, X., Willner, K., Zheng, Y., Damme, F., Bruzzone, L., Oberst, J., 2024. ELunarDTMNet: Efficient Reconstruction of High-Resolution Lunar DTM From Single-View Orbiter Images. *IEEE Transactions on Geoscience and Remote Sensing*, 62, 1-20.

Chen, H., Hu, X., Oberst, J., 2022. PIXEL-RESOLUTION DTM GENERATION FOR THE LUNAR SURFACE BASED ON A COMBINED DEEP LEARNING AND SHAPE-FROM-SHADING (SFS) APPROACH. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, V-3-2022, 511–516.

Creech, S., Guidi, J., Elburn, D., 2022. Artemis: An overview of nasa’s activities to return humans to the moon. *2022 IEEE Aerospace Conference (AERO)*, 1–7.

Falgout, J. T., Gordon, J., Lee, L., Williams, B., n.d. USGS Advanced Research Computing, USGS Hovenweep Supercomputer.

Goossens, S., Mazarico, E., Ishihara, Y., Archinal, B., Gaddis, L., 2020. Improving the geometry of Kaguya extended mission data through refined orbit determination using laser altimetry. *Icarus*, 336, 113454.

Group, C. S. M. W., 2007. Community Sensor Model (CSM) Technical Requirements Document (TRD). Technical Report, National Geospatial Intellegence Agency.

Haruyama, J., Hara, S., Hioki, K., Iwasaki, A., Morota, T., Ohtake, M., Matsunaga, T., Araki, H., Matsumoto, K., Ishihara, Y., Noda, H., Sasaki, S., Goossens, S., Iwata, T., 2012. Lunar Global Digital Terrain Model Dataset Produced from SELENE (Kaguya) Terrain Camera Stereo Observations. *43rd Annual Lunar and Planetary Science Conference*, Lunar and Planetary Science Conference, 1200.

Henriksen, M., Manheim, M., Burns, K., Seymour, P., Speyerer, E., Deran, A., Boyd, A., Howington-Kraus, E., Rosiek, M., Archinal, B., Robinson, M., 2017. Extracting accurate and precise topography from LROC narrow angle camera stereo observations. *Icarus*, 283, 122-137. Lunar Reconnaissance Orbiter - Part II.

Kato, M., Sasaki, S., Takizawa, Y., the Kaguya project team, 2010. The Kaguya Mission Overview. *Space Science Reviews*, 154(1), 3–19.

Kirk, R. L., Mayer, D. P., Fergason, R. L., Redding, B. L., Galuszka, D. M., Hare, T. M., Gwinner, K., 2021. Evaluating Stereo Digital Terrain Model Quality at Mars Rover Landing Sites with HRSC, CTX, and HiRISE Images. *Remote Sensing*, 13(17), 3511. Number: 17 Publisher: Multidisciplinary Digital Publishing Institute.

Kring, D. A., Bickel, V. T., van der Bogert, C. H., Fagan, A. L., Gaddis, L. R., Hiesinger, H., Hurtado, J. M., Joy, K. H., Lemelin, M., Looper, C. A., Osinski, G. R., Pösges, G., Siegler, M., Tikoo, S. M., Zacny, K., 2023. Elevation changes and slope that may affect eva workload near potential artemis landing sites. *2023 IEEE Aerospace Conference*, 1–17.

Kumari, N., Bretzfelder, J. M., Ganesh, I., Lang, A., Kring, D. A., 2022. Surface Conditions and Resource Accessibility at Potential Artemis Landing Sites 007 and 011. *The Planetary Science Journal*, 3(9), 224.

Laura, J. R., Archinal, B. A., Bland, M. T., Gaddis, L. R., Hagerty, J. J., Hare, T. M., Skinner, J.A., J., 2018. Planetary Spatial Data Infrastructure Foundational Data Product Knowledge Inventory. *LPI Contributions*.

Laura, J. R., Mapel, J., Hare, T., 2020. Planetary Sensor Models Interoperability Using the Community Sensor Model Specification. *Earth and Space Science*, 7(6), e2019EA000713. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1029/2019EA000713.

Liu, C., Du, W., Tian, X., 2017. Lunar dem super-resolution reconstruction via sparse representation. *2017 10th International Congress on Image and Signal Processing, BioMedical Engineering and Informatics (CISP-BMEI)*, 1–5.

Mapel, J. A., Berry, K., Paquette, A., Rodriguez, K., Stapleton, S., Hare, T. M., Laura, J. R., 2019. The Abstraction Layer for Ephemerides Library. LPI Editorial Board (ed.), *4th Planetary Data Workshop*, LPI Contributions, 2151, 7085.

Mazarico, E., Neumann, G., Smith, D., Zuber, M., Torrence, M., 2011. Illumination conditions of the lunar polar regions using LOLA topography. *Icarus*, 211(2), 1066-1081.

Nefian, A. V., Husmann, K., Broxton, M., To, V., Lundy, M., Hancher, M. D., 2009. A bayesian formulation for sub-pixel refinement in stereo orbital imagery. *2009 16th IEEE International Conference on Image Processing (ICIP)*, 2361–2364.

Pomerleau, F., Colas, F., Siegwart, R., Magnenat, S., 2013. Comparing ICP variants on real-world data sets. *Autonomous Robots*, 34(3), 133–148. https://doi.org/10.1007/s10514-013-9327-2.

Smith, D. E., Zuber, M. T., Neumann, G. A., Mazarico, E., Head, J., Torrence, M. H., Lola Science Team, 2011. Results from the Lunar Orbiter Laser Altimeter (LOLA): Global, High Resolution Topographic Mapping of the Moon. 2350. Conference Name: 42nd Annual Lunar and Planetary Science Conference ADS Bibcode: 2011LPI....42.2350S.

Speyerer, E. J., Robinson, M. S., Bussey, B., 2025. Optimizing Lunar Exploration: A Multidimensional Approach to Landing Site Selection. *IEEE Geoscience and Remote Sensing Letters*, 22, 1-5.

Tooley, C. R., Houghton, M. B., Saylor, R. S., Peddie, C., Everett, D. F., Baker, C. L., Safdie, K. N., 2010. Lunar Reconnaissance Orbiter Mission and Spacecraft Design. *Space Science Reviews*, 150(1), 23–62.

Wu, B., Dong, J., Wang, Y., Rao, W., Sun, Z., Li, Z., Tan, Z., Chen, Z., Wang, C., Liu, W. C., Chen, L., Zhu, J., Li, H., 2022. Landing Site Selection and Characterization of Tianwen-1 (Zhurong Rover) on Mars. *Journal of Geophysical Research: Planets*, 127(4), e2021JE007137. e2021JE007137 2021JE007137.

Yokota, Y., Haruyama, J., Honda, C., Morota, T., Ohtake, M., Kawasaki, H., Hara, S., Hioki, K., 2008. Lunar topography: Statistical analysis of roughness on a kilometer scale. *Advances in Space Research*, 42(2), 259-266.

Zhong, Z., Yan, J., He, H., Wen, Q., Liu, D., Barriot, J.-P., 2023. Illumination and regolith temperature at China’s next candidate lunar landing site Shackleton crater. *Science China Earth Sciences*, 66(2), 417–429.