



RESEARCH ARTICLE

10.1029/2025EA004594

Special Collection:

The Lunar Trailblazer Mission
Collection: Mission, Instruments,
Data Analysis Plan

Post-Acquisition Image-Based Localization for High Resolution Thermal and Visible/Shortwave Infrared Images With Application to the Lunar Trailblazer Mission

Kevin D. Gauld¹ , James L. Dickson² , Tristram J. Warren³ , Jihoon Yang^{1,4} , Neil Bowles³ , and Bethany L. Ehlmann¹ 

¹California Institute of Technology, Pasadena, CA, USA, ²University of Minnesota, Minneapolis, MN, USA, ³University of Oxford, Oxford, UK, ⁴University of California, Berkeley, CA, USA

Key Points:

- We present a process for post-acquisition, image-based data localization for satellite missions with use for the Lunar Trailblazer mission
- We demonstrate image matching success across short- and mid-wave infrared image wavelength bands and topographic hillshade data sets
- This IMPPAIL procedure is applicable to future missions that require higher pointing accuracy than is possible with hardware solutions

Correspondence to:

K. D. Gauld,
kgauld@mit.edu

Citation:

Gauld, K. D., Dickson, J. L., Warren, T. J., Yang, J., Bowles, N., & Ehlmann, B. L. (2026). Post-acquisition image-based localization for high resolution thermal and visible/shortwave infrared images with application to the lunar trailblazer mission. *Earth and Space Science*, 13, e2025EA004594. <https://doi.org/10.1029/2025EA004594>

Received 2 JUL 2025

Accepted 23 FEB 2026

Author Contributions:

Conceptualization: Kevin D. Gauld, James L. Dickson, Jihoon Yang, Bethany L. Ehlmann

Data curation: Kevin D. Gauld, James L. Dickson, Tristram J. Warren

Formal analysis: Kevin D. Gauld

Funding acquisition: Neil Bowles, Bethany L. Ehlmann

Investigation: Kevin D. Gauld, James L. Dickson, Bethany L. Ehlmann

Methodology: Kevin D. Gauld, James L. Dickson, Jihoon Yang, Bethany L. Ehlmann

© 2026. The Author(s).

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial License](https://creativecommons.org/licenses/by-nc/4.0/), which permits use, distribution and reproduction in any medium, provided the original work is properly cited and is not used for commercial purposes.

Abstract The Lunar Trailblazer mission aimed to assess the presence of water on the lunar surface using imaging spectroscopy in visible shortwave infrared (VSWIR) coupled with high-resolution multispectral imaging in thermal midwave-infrared (MWIR), captured simultaneously over the same target from orbit around the Moon with two different instruments. Uncertainties in clock timing, instrument models, and instrument pointing knowledge manifest as geospatial offsets between the two data sets that must be corrected in post-processing to enable co-registration, tying the acquired images to their precise latitudes and longitudes on the Moon. This work describes an algorithmic approach to co-registering and geolocating images after acquisition without high precision instrument and spacecraft pointing models, the Iterative Matching Pipeline for Post-Acquisition Image Localization (IMPPAIL), utilizing previously acquired data for development. We use Lunar Orbiter Laser Altimetry (LOLA) and Kaguya data to make shaded relief maps as the basemap on which to project data. To test our processing pipeline prior to Lunar Trailblazer data collection, we use Moon Mineralogy Mapper (M³) data for VSWIR images and simulated MWIR images. When demonstrated on these data sets, IMPPAIL produces a 98% success rate registering VSWIR data to LOLA/Kaguya shaded relief maps and successfully co-registered MWIR and VSWIR in all four simulation cases. We include a code package with software tools allowing this algorithm to be used for a variety of data sets across many other missions.

Plain Language Summary The Lunar Trailblazer spacecraft was designed to map surface temperature and water content on the Moon via satellite imagery. Onboard are two imagers that measure thermal and visible light bands, which must be aligned to be used in conjunction with each other as well as located precisely relative to pre-existing topography data. To do this, we develop an algorithm which iteratively matches key features between the images, allowing for the co-registration of data points. We test the algorithm on topography and visible imagery data from a past satellite mission and simulated data for thermal images. We show 98% accuracy and success in both imagery wavelength bands. While this was originally created for use on the Lunar Trailblazer mission, the process is applicable to future missions requiring similar functionality, and we make the code available for other users.

1. Introduction

Water ice has been reported on the Moon (Colaprete et al., 2010; Feldman et al., 1998; Li et al., 2018; Rubanenko et al., 2019; Spudis et al., 2013), as well as non-polar H₂O or OH (Clark, 2009; Li & Milliken, 2017; Pieters et al., 2009; Sunshine et al., 2009), inspiring exploration of the distribution and abundance of water and other volatiles on the lunar surface. The Lunar Trailblazer mission aims to map volatiles on the lunar surface and understand the lunar water cycle (Ehlmann, Klima, et al., 2026; Ehlmann et al., 2022 this collection). Imaging the surface from a 100 ± 30 km orbit, Lunar Trailblazer uses the High-resolution Volatiles and Minerals Moon Mapper (HVM³; Bender et al., 2022; Thompson et al., 2026) and Lunar Thermal Mapper (LTM; Bowles et al., 2026) instruments to record images in visible and shortwave infrared (VSWIR) and thermal mid-infrared wavelengths, respectively. These imagers provide information at <100 m/pixel scale (nominally ~70 m/pix and 55 m/pix from 100 km, respectively, see Ehlmann, Klima, et al., 2026) on the volatiles, composition, and temperature on the surface of the Moon, acquiring data of the same location simultaneously from orbit.

After both images are collected, they must be positioned on the lunar surface to be analyzed in conjunction with each other and prior lunar data sets. Image acquisition is sequenced utilizing a target center latitude and longitude,

Project administration: James L. Dickson, Neil Bowles, Bethany L. Ehlmann
Resources: Kevin D. Gauld, Neil Bowles, Bethany L. Ehlmann
Software: Kevin D. Gauld, James L. Dickson, Tristam J. Warren, Jihoon Yang, Neil Bowles
Supervision: James L. Dickson, Neil Bowles, Bethany L. Ehlmann
Validation: Kevin D. Gauld, James L. Dickson, Bethany L. Ehlmann
Visualization: Kevin D. Gauld
Writing – original draft: Kevin D. Gauld, Bethany L. Ehlmann
Writing – review & editing: Kevin D. Gauld, James L. Dickson, Tristam J. Warren, Neil Bowles, Bethany L. Ehlmann

so the general locations of data are known. Uncertainties in pointing, positioning knowledge of the satellite, instrument mounting accuracy, and instrument camera models limit the precision on the location of recorded images. Furthermore, differences in the pointing of the separately mounted instruments and the different pixel sizes of the two imagers create a discrepancy in the portions of the surface observed between the thermal radiance and VSWIR radiance images. As such, Lunar Trailblazer requires a method to register the two acquired images to each other, and then register the data to the surface of the Moon. This post-processing ensures accurate localization of images and allows Lunar Trailblazer data to be used with other existing high-resolution lunar data sets, including photometric correction of Trailblazer data using angles derived from lunar topography in later processing steps.

Some instruments and missions approach this problem with precision timing and navigation coupled with detailed camera models (e.g., HiRISE, McEwen et al., 2007; LROC, Robinson et al., 2010). Yet, there typically remain offsets of several pixels between data sets. As a low-cost planetary mission, Lunar Trailblazer had to make design trades to fit in the program's cost constraints (see Ehlmann, Klima, et al., 2026; Ehlmann et al., 2022 for further discussion), accepting some limits and adapting requirements and approaches. Lunar Trailblazer's spacecraft clock and guidance, navigation, and control hardware provide information that, coupled with Trailblazer's orbital velocity, can lead to image position uncertainties greater than those of higher class missions, up to a few kilometers (see Section 2.2). This makes post-acquisition co-registration mandatory for analysis of both data sets together given sub-100 m/pix resolution. Even in the lower uncertainty cases for higher class missions, misregistrations can impede the analysis of overlapping data at native resolution. Consequently, automated approaches to improving image co-registration are desirable.

Furthermore, the science requirements for Lunar Trailblazer's image acquisition do not have strict absolute pointing knowledge requirements. Acquisition is required to include the center latitude and longitude targeted and a 3-km × 3-km buffer around for geologic context, amply met given both instruments' swath widths of >~20 km. However, after images are acquired, for co-registration with other lunar data sets, approximately single pixel-scale accuracy in locations relative to lunar topographic base maps is required for accurate projection of the data, including orthorectification, and is part of determining incidence and emissions angles for a full photometric correction. This also provides a reference coordinate frame on which to compare and co-analyze data from Lunar Trailblazer with that of past imaging missions.

Because of decades of exploration of the Moon, we are fortunate to have existing high resolution topography and imagery. Consequently, the science team decided to invest in a pipeline procedure for after-acquisition localization and co-registration of images. Similar approaches have been used for large-scale mosaicking projects on Mars (e.g., Dickson et al., 2024; Edwards et al., 2011) and for correcting the locations of Moon Mineralogy Mapper (M³) data (e.g., Boardman et al., 2011). Here, we describe the approach that the Lunar Trailblazer mission developed to localize HVM³ and LTM images on the surface. We demonstrate the use of the Scale-Invariant Feature Transform (SIFT) algorithm (Lowe, 1999) to register imagery data to a Lunar Orbiter Laser Altimeter (LOLA)/Kaguya topographic map, the reference data set for localization on the lunar surface at low- and mid-latitudes, and LOLA alone at high-latitudes. We show how, first, HVM³-like and LTM-like data are co-registered and then how both are referenced to LOLA to provide accurately georeferenced products for mapping volatiles, composition, and thermophysical properties. After this manuscript was submitted, the Lunar Trailblazer mission was lost (Ehlmann, Bellerose, et al., 2026); for simplicity and potential reuse if reflown, we retain the description of the workflows as-designed.

2. Methods

2.1. Workflow Overview

Matching LTB data to the lunar surface requires a benchmark reference data set for the lunar surface with known latitude and longitude. The Lunar Orbiter Laser Altimeter (LOLA)/Kaguya SLDEM 2015 merged DEM provides a map of the lunar topography at a resolution of 60 m/pixel within 60° of the equator, combined with LOLA data rendered at 60 m/pixel poleward of 60° in each hemisphere, providing comparable resolution and accurate features to register collected images (Barker et al., 2021; Smith et al., 2010; Zuber et al., 2010). (Note that higher resolution LOLA data sets became available during the course of this work, but we proved the workflow with the lower resolution data.)

Table 1

Key Parameters for Data Expected From HVM³ and the Corresponding Values Used to Simulate Those Parameters Using M³ and LOLA/Kaguya

Parameter	HVM ³ data	M ³ mimicking HVM ³	LOLA/Kaguya hillshade
Spectral Range	600–3,600 nm	450–3,000 nm	N/A (hillshade)
Resolution (m/pix)	50–90 m/pix (depending on satellite altitude) 70 m/pix nominal average	~70–200 m/pix (depending on satellite altitude) 140 m/pix nominal average	60 m/pix in this study (can use <30 m/pix resolution with HVM ³ in flight)
Nominal Image size	22 km cross-track 27 km along-track	42.6 km cross-track 52.1 km along-track	64.4 km cross-track 112.7 km along-track (based on acquisition uncertainty; Section 2.3.2)

Note. The LOLA/Kaguya hillshade will be generated via the same processing based on the acquisition uncertainty in cross-track and along-track pointing relative to the image center point, but the output size of the hillshade will change when using HVM³ versus M³.

Once data are returned from Lunar Trailblazer, we will register the LTM image to HVM³, then register the HVM³ image to a projected LOLA/Kaguya hillshade image, subset in the vicinity of the target and computed for the viewing geometry and local solar time of the observation. By doing this stepwise, we align thermal infrared features seen in the LTM image directly to visible/shortwave infrared features seen by HVM³. This has two benefits: (a) it allows co-analysis of the two data sets to within ~1 pixel in “detector space,” that is, in the surface data as recorded by discrete detector elements, and (b) it requires only one image transformation for each to be projected. The first is beneficial because instrument calibration approaches (and artifact removal steps) are often row and column dependent. Considering relative radiance variation by channel and spatial position is best done before blending of pixels or interpolation/resampling. This allows measured temperature and water content to be directly compared without blurring relative to the inherent scale of measurement. The second is important as a best-practice approach of preserving discrete spatial sampling of surface information to the greatest degree possible through the creation of data products. Unprojected Lunar Trailblazer images are inherently higher fidelity than projected images, which resample and blend adjacent pixels in an image transform operation.

2.2. Data Set Preparation

As no Lunar Trailblazer data are presently available, the pipeline was simulated for testing using a combination of existing data sets and models.

2.2.1. HVM³-Simulating Data Preparation With M³

To mimic HVM³, we use data from the Moon Mineralogy Mapper (M³) imager, which was flown on Chandrayaan-1, collecting data from 2008 to 2009 (Boardman et al., 2011; Green et al., 2011). M³ collected radiance images of the Moon that are similar in structure to those which will be captured by HVM³ from Lunar Trailblazer. As such, images from M³ are used to validate the localization procedure, simulating HVM³. Also, as M³ contains most of the same spectral bands we will see in HVM³ (Table 1), the use of this data provides confidence that successes will translate to data collected by HVM³. We used 404 images from between 19 November 2008 to 15 August 2009, within 80° to –80° latitude, acquired in mapping mode (140 m/pixel), corresponding to optical periods 1A-2B, as the test data set because this operational period is the most representative of the observation conditions we expect from Lunar Trailblazer during the majority of observations.

While the M³ data is similar, we preprocess the data set to ensure it matches HVM³ scalings as closely as possible in spatial resolution relative to LOLA and spatial extent (Table 1). The M³ image is cropped to a size consistent with the HVM³ image aspect ratio. HVM³'s cross-track size is set by its cross-track swath width, dictated by hardware and data downlink. It is 22 km in typical “half-width” mode with capability to acquire 44 km full width images in special “full-width” mode. The along-track HVM³ image size is set by mission system-level imaging requirements set to account for pointing uncertainty. The system requirement is that HVM³ must image a target surface coordinate with a 3 km surrounding radius within a 10 × 10 km square region. Including all potential errors from the clock, orientation, and other sources, this 10 × 10 km required region must account for pointing error spanning about 11 km along track and 2 km cross track (Figure 1). Thus, HVM³'s typical image acquisition

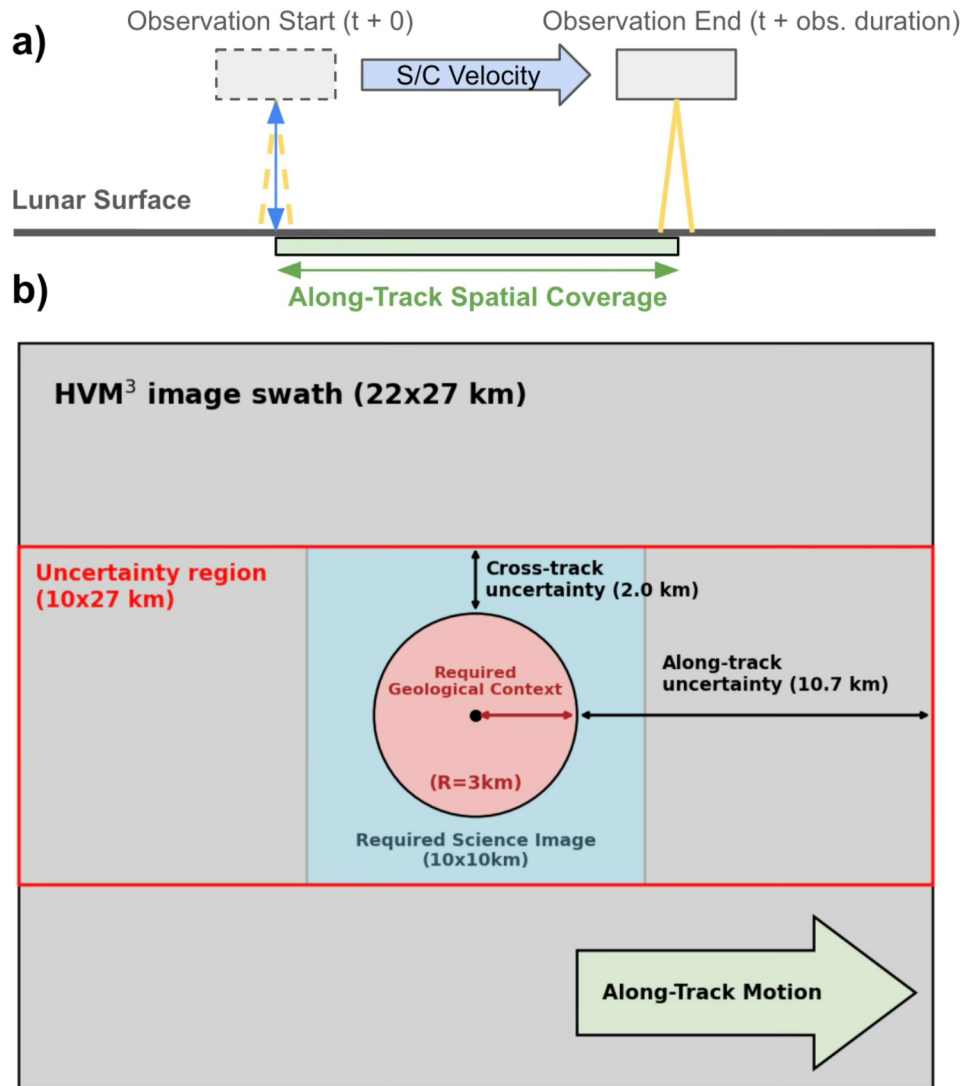


Figure 1. (a) Diagram of satellite motion as images are collected via pushbrooming across the surface. (b) diagram of the HVM³ image swath region spanning from the observation start to observation end, with the required minimum science image in the center, surrounded by uncertainty. Sources of uncertainty include instrument co-alignment, pointing, clock, control sync, and ephemeris uncertainties. At 100 km orbit, the mapping error on the central point can be up to 2.0 km cross-track and 10.7 km along-track.

at 100 km is 22 km cross-track \times 27 km along track (320 pixels cross-track \times approximately 400 pix along track) at a resolution of approximately 70 m/pix.

Because the M³ data resolution is typically 140 m/pix, we scale the image size to match the number of pixels which will be used in the matching process at higher resolutions. Scaling by a factor of 2.3x, we arrive at a target dimension of 51.3 \times 63 km (360 \times 450 pixels) for testing. However, for testing we are limited by the dimensions of the M³ imager, which only has 304 cross-track pixels, so our dimensions for testing are instead 304 \times 372 pixels, corresponding to a swath of 42.6 \times 52.1 km. M³ includes 85 bands from 450 to 3,000 nm. To increase signal/noise and thus improve the probability of a successful match, we averaged 10 neighboring bands in I/F over 2.4–2.8 μ m and used this longest wavelength data as the M³ imagery data used for testing the matching process.

2.2.2. LTM-Simulating Data Preparation

While M^3 provides an analog to the HVM^3 imager, sourcing imager data in MWIR (7–10 μm) is challenging, as past missions (e.g., Diviner; Paige et al., 2010) have collected thermal data at coarser resolutions (~ 200 m/pix compared to our 55 m/pix) and not in pushbroom format. Also, we require thermal test data within the image regions captured by M^3 , restricting the collected data we are able to use, as we register the MWIR image directly to the captured VSWIR/ M^3 image rather than to ground topography to ensure captured data is registered in detector space, independently of its transform onto the ground data set.

Our solution to this is to simulate expected MWIR data seen by LTM within a given region for many different lunar times. Given the long computational time required to create each simulated thermal image, we choose a few representative data points to spot check performance of the matching algorithm. VSWIR and MWIR produce images due to fundamentally different phenomena: mostly reflected and mostly thermally emitted light, respectively. Four image samples exhibiting various combinations of these features were chosen to test the thermal model. The spot check site data points are chosen for their number of “features” (many or not many craters), and size of features (large or small craters).

We use a LOLA DEM to simulate a LTM thermal image with the expected 55 m/pix resolution, incorporating the LTM instrument's thermal response and noise profiles. The thermal model used is described in King et al. (2020) and Hayne et al. (2017). Briefly, ray tracing is applied using a LOLA (SLDEM2015; Barker et al., 2016) DEM as input to calculate local shadowing and solar input energy, considering local slope and topography. The one-dimensional thermal diffusion equation is solved at each time step to predict temperatures at multiple subsurface depths. The modeled surface temperature is then converted into radiance, and the calculated radiance is convolved with the LTM's instrument point spread function (PSF). Instrument noise is added to simulate the expected output from the instrument. This process is repeated for each of the LTM's spectral channels, as each channel has a different bandpass and PSF.

Because HVM^3 and LTM will have similar resolutions and sizes, for simplicity in our trials, we ensure the simulated thermal image has the same resolution as the captured M^3 image, which is 140 m/pix. To achieve this resolution, we downsample using a pixel area relation, averaging windows of pixels to bring the simulated resolution up to 140 m/pix. The along-track dimension is a 9° FOV for LTM and a 12° FOV for HVM^3 (Bowles et al., 2026; Ehlmann, Klima, et al., 2026). As such, we crop the simulated LTM image such that it occupies 80% of the width of the M^3 image (cross-track) and the full along track distance, producing a swath of 228×372 pix, matched with an M^3 image with dimensions of 304×372 pix.

2.2.3. Hillshade Data Preparation

The benchmark data set is the LOLA/Kaguya merged DEM at 60 m/pix (just LOLA poleward of 60°) (Barker et al., 2016), which we use to create a shaded relief map based on the observing conditions when the corresponding M^3 image was captured. By matching surface illumination, including the effects of topography, in the locale of the M^3 image, we create a shaded relief image with features closely matching those observed in the collected image.

To create the hillshade for any given region of the Moon, the LOLA/Kaguya DEM is first cropped by latitude/longitude on a cylindrical projection of the lunar surface around the targeted centerpoint. With an expected image swath of 22×27 km and uncertainties of $\sim 6 \times 22$ km, the hillshade should be 28×49 km for the HVM^3 case. For testing with M^3 , we scale by the relative pixel scaling of 2.3x, providing a hillshade area of 64.4×112.7 km ($1,284 \times 2,250$ pix). This size is chosen to ensure the image falls within the hillshade region while minimizing the area outside of the maximum uncertainty to limit the number of false matches. This hillshade size is not a hard constraint, and we test larger hillshade sizes for sensitivity (Section 3.5). We form the hillshare region by transforming the center lat/lon to x/y surface coordinates, creating a boundary box in x/y, then transforming the boundary back into lat/lon coordinates. This boundary is then used to crop the DEM, defining the image region used in matching. After retrieving the cropped DEM, it is reprojected using a sinusoidal projection about the center longitude, ensuring the areas represented in the hillshade are similar in size to those in the captured image. After this reprojected, the shaded relief is generated using the GDALDEM command line tool using the solar incidence and azimuth at the local solar time of observation. This simulates the shadows present when imaged by




Figure 2. Intermediate products used to generate a final LOLA hillshade relief image (right). First, the LOLA/Kaguya DEM is cropped to a lat/lon region (left), after which it is reprojected about its center using a sinusoidal projection (center). The shaded relief is created with the GDALDEM command line tool with observation parameters matching those of the observed M^3 image.

HVM³. The 8-bit raster hillshade image centered around the science target will be used for localizing the HVM³ image. (Figure 2).

2.3. Iterative Feature Matching

We create an iterative matching pipeline for post-acquisition image localization (IMPPAIL). To match the two images, we detect common features between the two images and generate match pairs between SIFT key-points. Then a subset of common features of greatest confidence is chosen, discarding other potential key-points. Once multiple match pairs are created, they are used to compute a transform between the two images. The computed transform is then applied to the input image, and the matching repeats on this transformed image, repeating up to 10 times to ensure a high-quality set of match pairs (Figure 3).

In the first iteration of the process, the applied transformation is a “best guess,” where the M^3 image is scaled to match its expected extent in the hillshade based on the relative ground resolutions, and is centered in the hillshade frame. In subsequent iterations, the homography computed by the prior iteration is used as the initial transformation, as it is the current best-guess estimate of the overall homography.

2.3.1. Feature Matching: SIFT

A series of computer vision algorithms were tested for effectiveness and efficiency with the simulated data used in this study, and SIFT, an algorithm to detect multi-scale features in an image (see Lowe, 1999 for details), performed best. This result is consistent with prior work from the M^3 mission which successfully applied SIFT to ground-register M^3 data to lunar topography (Boardman et al., 2011). Briefly, SIFT can detect regions of high local contrast and then create key-points at the center of these features along with a descriptor encoding local image gradients, allowing for feature-matching that accommodates scale, rotation, and illumination changes. SIFT generates key-points in both a “source” and “destination” image, which are then matched together using a Brute-Force Matcher, which ties matching key-points across images based on the descriptor characteristics. These match-pairs ultimately allow for the creation of mappings which bring pixels from the source image (in column/row space) onto corresponding pixels in the destination image (in geographic lat/lon space).

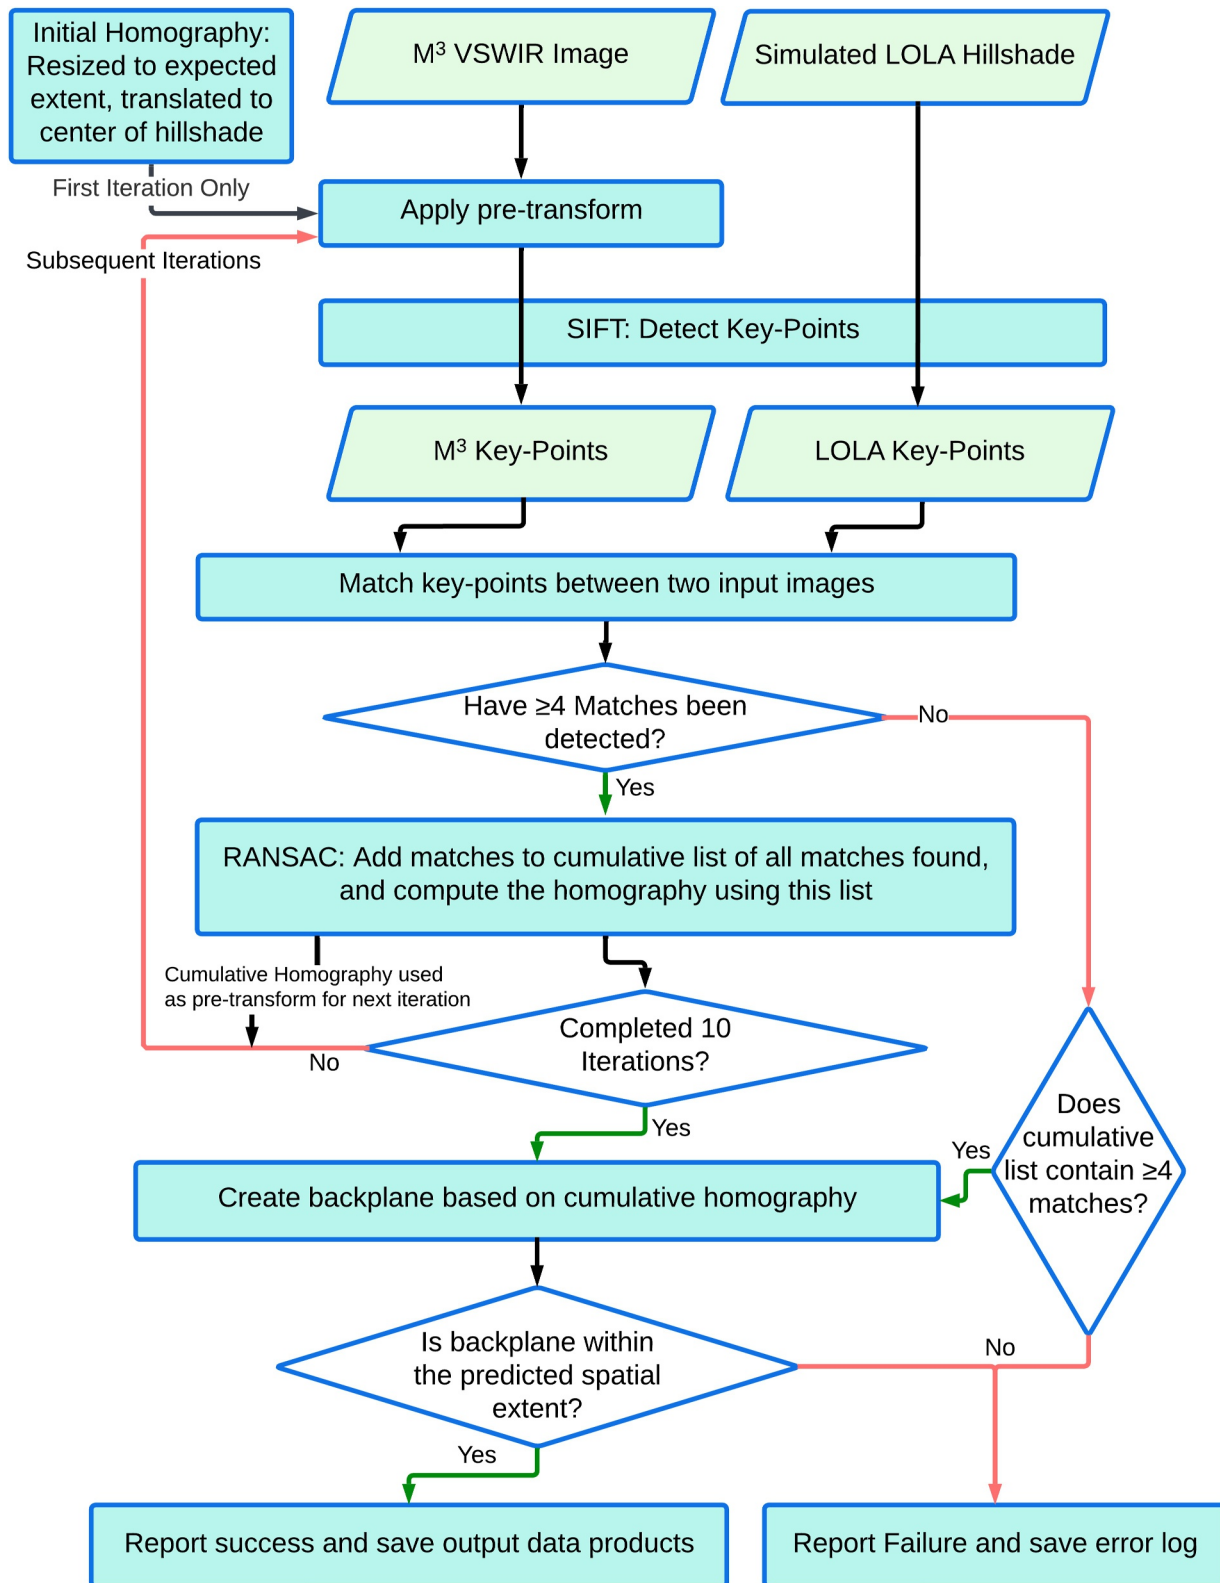


Figure 3. Flowchart of the IMPPAIL processing described in Sections 2.3 and 2.4. For thermal data, the simulated LTM images replace the M³ images, and the LOLA hillshade is replaced by the M³ image data. No other changes are necessary, as the algorithm operates the same given a different input data set.

Table 2
Parameters Used for Our Implementations of SIFT and Brute-Force Matching, Alongside Descriptions

SIFT parameter	Description	Default value	Our value
nFeatures	The number of best features to retain, ranked by local contrast. 0 to preserve all features.	0	0
nOctaveLayers	The number of layers in each octave when finding local contrast maxima.	3	4
contrastThreshold	The lower-bound on the contrast to filter out weak features. The larger the threshold, the less features are produced by the detector.	0.04	0.03
edgeThreshold	The upper-bound on the contrast to filter out edge-like features. The larger the threshold, the more features are retained by the detector.	10	15
sigma	The sigma of the Gaussian applied to the input image during the first octave. The sharper the image, the higher the sigma should be.	1.6	1.6
Matching parameter	Description	Default value	Our value
normType	Norm to use for computing distance between two key-point descriptors	L^2	L^2
crossCheck	Flag to indicate whether we require bidirectionality when creating match-pairs	False	False
k	Number of nearest neighbor matches to find, sorted by descriptor distance	1 (for only best match)	2
r	Ratio used in Lowe's Test, thresholding the acceptable uncertainty between the two nearest neighbor distances	N/A (not used by default)	0.8

Note. These are largely based on the parameters used in Lowe (1999); Lowe (2004). See the code supplement (Gauld et al., 2026a, GitHub/Zenodo) for the implementation of these parameters. The last two matching parameters have no formal default values, as they are not included in the matcher by default and are implemented as a filter in-post. These parameters are application dependent and require tuning for use with other data sets.

We use the standard feature detector and matcher used in most implementations of SIFT (see code supplement), which uses a difference of gaussians (DoG) in the local pixel region to find contrast gradients. SIFT constructs feature representations (key-points) iteratively through “octaves,” where each octave consists of progressive Gaussian-blurred image “layers,” followed by a downsample of factor 2, which signifies the beginning of the following octave. The DoG is computed for each layer and its predecessor, with any local extrema in these images marked as key-points. After bounding the gradients to ensure low contrast noise and high contrast edges are not counted as features, SIFT provides the detected key-points. The algorithm parameters able to be modified are the number of layers in each octave, the number of octaves, the thresholds on feature detection, and the sigma factor, that is, the width, for each Gaussian blur (Table 2). For our implementation, we ensure that many match-pairs are created, widening the thresholds for match creation and increasing the number of layers per octave. This ensures that enough correct correspondences are detected for homography estimation, even though it also increases the number of outliers. We opt to maximize this match generation to retain match-pairs and rely on subsequent use of the RANSAC algorithm (see Section 2.3.2) to isolate the true inliers, computing the true homography.

After detecting key-points in each image using SIFT, we match them across the two images. Each key-point has a descriptor characterizing the local contrast gradient, which is compared in the matching process. To compute distance between descriptors, we use the Euclidean (L^2) norm. We also disable cross-checking, which checks pairs in both directions to assure that if keypoint A has a nearest-distance keypoint B, then a match is only created if B has A as its nearest-distance pair. We found that it is too aggressive of a filter for our application, removing too many true matched pairs to find the correct homography. Rather than cross-checking, to improve reliability, we discard matches which have high ambiguity in the matching routine. To accomplish this, we use a K-Nearest-Neighbors approach with $k = 2$ to find the two closest matches for the given source key-point. We then apply Lowe's ratio test (Lowe, 2004), taking the ratio of the closest distance over the second-closest distance. If this exceeds a ratio of 0.8, the match is discarded due to ambiguity in determining its closest match pairing. While this can filter out some valid matches and include some invalid matches, this ratio score performs sufficiently well in the vast majority of cases, enabling each image to have an adequate number of pairs to proceed to the next steps.

2.3.2. Homography Calculation

Now that pairs have been detected, we use Random-Sample Consensus (RANSAC; Fischler & Bolles, 1981) to find a function mapping the key-points in the source image with their matches in the destination image. This requires the creation of a homography, which is a transformation to bring the pixel coordinates of one image onto




Figure 4. Tie-points after post-processing with RANSAC. The left side of the image is the M^3 image, and the right side is the corresponding hillshade. In the initial matching step, all of the matches are found and marked as key-point matches (red). During the RANSAC step a consensus is found (highlighted in green) which is used to compute the homography. Red pairs designate erroneous matches disobeying the RANSAC consensus and are thus not used in the mapping to bring M^3 onto the hillshade.

the pixels in another image. This mapping can only be computed with a minimum of four pairs, as the mapping is not well-defined for fewer inputs. RANSAC use ensures the computed homography filters out any incorrect matches, preserving matches which form a consensus. RANSAC works by randomly sub-sampling the matches, creating a homography based off of the sub-sample. After creating many random sub-sample matches, these are evaluated on every match pair in the image. The homography with the most match-pairs in agreement with the homography (inliers) is chosen as the “consensus,” then the homography is recomputed using only inlier points, ensuring that all erroneous outlier pairs are filtered out, and that the homography is accurate (Figure 4). Valid pairs outside the consensus have no bearing on the match quality, as RANSAC ensures there are enough points within the consensus to generate a high-quality homography. If no consensus is found, the homography is recalculated using a least-squares approximation on all match-pairs, which is more susceptible to the effects of outliers. The resultant homography is a projective transformation, enabling correct alignment by accounting for any perspective misalignments between the hillshade and captured image. Given the relatively small image swath, the surface can be locally approximated as planar, making a homography an appropriate transformation model. Upon computing this transform, the accuracy is not yet quantified, as this is done in the last step of the algorithm once backplanes are created. As such, both RANSAC and least-squares may fail in this step, resulting in an incorrect homography which is later filtered out.

In our case, homographies are calculated twice: once registering LTM to HVM^3 , and once registering HVM^3 to the LOLA/Kaguya hillshade. In both cases, RANSAC is extremely useful, as the SIFT implementation is prone to creating false match pairs, though these pairs often only account for a minority of the detected matches. By using RANSAC, these incorrect matches are neglected, allowing for more lenient bounds on the key-point creation (see Section 2.3.1). By ensuring wide bounds on key-point generation, we optimize the number of correctly detected match pairs and minimize the number of falsely rejected pairs.

2.3.3. Iterative Feature Matching

We use the SIFT and RANSAC procedures iteratively to improve the matching capability of IMPPAIL. After running through one iteration of SIFT, we are able to use RANSAC to produce a homography registering the hillshade and observed image. Now, we feed this homography back into the input stage of SIFT, transforming the M^3 input to more closely match the expected spatial properties of the LOLA hillshade. The process then repeats. Each iteration, we update a running list of all match-pairs used from the past iterations, and use this cumulative list to compute the initial input transformation. We choose to repeat this process for up to 10 iterations, stopping at the first iteration where fewer than four new match-pairs are detected, with four pairs being the minimum requirement to create a valid homography. At the end of the iterative process, the output match-pairs are used to create a stable homography which co-registers the two images at a high resolution or reports failure (Section 2.4.2).

2.4. Output Products

2.4.1. Location Backplanes

After the final homography is calculated, we use it to create latitude and longitude backplanes to map each pixel in the original image to a coordinate. We begin by creating a mask image to select a region about 10 pixels larger in each dimension than the original M^3 image to prevent destruction of boundary data points. This mask region is transformed into the frame of the hillshade using the homography found during feature matching. This transformed mask is then used to create a gridded raster file to populate latitude and longitude from the known values in the LOLA hillshade. Then we perform the inverse transform to bring the grid region back into the reference of the original image. Once the extra 10 pixels are trimmed, we have a populated latitude and longitude grid for each pixel in the original acquired image.

2.4.2. Measuring Successes

We must also define what success and failure are in the context of this matching process. For the feature matching step, a homography cannot be computed if fewer than 4 match-pairs are found. In this case, there is no backplane generated, and the match fails. This failure mode only occurs during the first SIFT iteration, as the algorithm only continues to the second iteration if 4 or more match-pairs are found in the first iteration. If there are 4 or more match-pairs computed, the homography can still be incorrect. From observations of these failure cases, most fail very badly, transforming the image area very far outside of the predicted size. We know the physical size of an HVM^3 image on the surface given the orbit parameters (Figure 1), so we filter based on expected size. If the match region is much larger or smaller than the expected swath of the recorded image, we mark it as a failure.

To ensure this approach accurately classifies failures and successes, we manually verified each of the images in our data set and found that this filter does accurately sort successes from failures. This manual verification showed that all images classified as a success were true successes, though some successes were erroneously marked as failures due to differences in uncertainty bounds (see Section 3.2). For Lunar Trailblazer flight use, the manual verification will be useful but not required, as the classification of successes/failures is consistently accurate. Upon both manual and automated inspection, the image overlays the hillshade region accurately to near pixel-scale, providing precise localization (see Section 3.4). In the future, the filter could also be expanded to ensure the backplanes are not discontinuous within the region and have consistent gradients from left to right and top to bottom, though this was not necessary to accurately classify the failures observed in our data set.

3. Results

3.1. Successes and Failures

We measure the success or failure based on the accuracy of the homography area in localizing the image region on the surface. For cases in which the matching works, the homography correctly places the area within the desired region, while incorrect transformations result in the images with overlaid regions different from our expected area with incorrect area scaling (Figure 5). To check the amount of misalignment, we use the average pixel offset across all match pairs (see Section 3.2). For these success cases, the match accuracy is within 1–2 pixels of misalignment (Section 3.4), and for failure cases there is always no alignment, so the classification is binary.

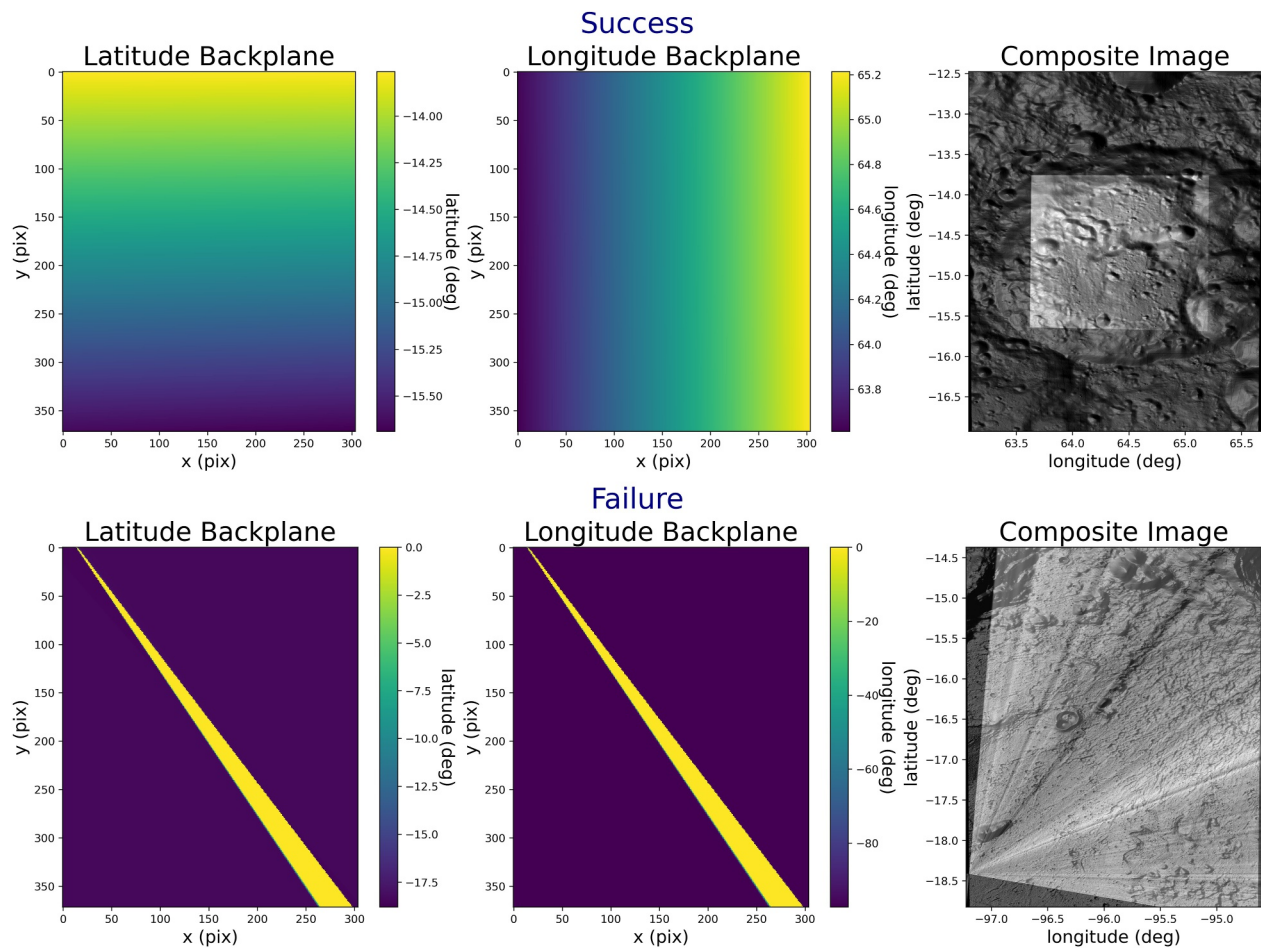


Figure 5. Example of a success (top) and failure (bottom) with their resultant backplanes. For the failure case, the range in both latitude and longitude is at least an order of magnitude higher than the range in the success, a clear indication of failure. Another notable indication is the sharp edges present in the backplane rather than a uniformly smooth gradient, though this is not explicitly checked in our processing. The lat/lon shown in the composite image is the ground truth lat/lon from the LOLA hillshade crop.

3.2. Aggregate Matching Results—VSWIR HVM³/M³ and Hillshade

In general, IMPPAIL works well to localize M³ images onto the lunar surface, reporting success in 388 of the 404 test cases. In reality, 396 of 404 are localized, a 98% match rate. When visually inspecting the data, we observe 8 images at high latitudes were incorrectly flagged as failures due to the correct match region being only partially contained in the bounds of the hillshade region, causing the backplane to not populate coordinates for the image region outside of the hillshade. (Figure 6). This is not due to the IMPPAIL algorithm but rather the fact that the pointing uncertainty of M³ is actually greater than that expected of HVM³, in some cases up to tens of kilometers in each direction (Boardman et al., 2011). As such, by restricting the hillshade to our expectation of HVM³'s uncertainty, we cause false-failures where the uncertainty in M³ operation is greater than we expect from HVM³'s acquisition. We consider these points to be successful, as they demonstrate the algorithm working as expected, correctly matching related features to each other across the images.

In looking at other potential sources of true failures, we find that some points at very high or low solar incidence angles account for most failure cases (Figure 7). Two failures occur for high-latitude images at high incidence angles. For these two images, the failure likely results from amplified shadows, where features “smear” together as shadows grow very large.

Three images fail when the solar angle is <30° and shadows are few to nonexistent, resulting in a breakdown of the matching algorithm, which heavily relies on contrast. Our results are consistent with those from Boardman et al. (2011) who found that M³ data with low solar zenith angles required manual registration to LOLA. This




Figure 6. Examples of each match condition, where M^3 image data is overlaid onto the LOLA hillshade. The left side of the center case exits the image frame, causing the algorithm to report a failure despite the match being successful.

appears partially related to the details of the hillshade creation function. Our process currently uses the GDALDEM (GDAL/OGR Contributors, 2023) command line tool, which creates a shaded relief model for the region and observation conditions provided. This estimate generally works very well but as the sun approaches nadir, the processing begins to fail, as shadows are not well-generated in contrast scaling within the algorithm, that is, the 0–255 dynamic range is not used in the product. If the initial matching procedure fails, we re-scale the elevations used by GDALDEM with the zFactor parameter. This scale factor is applied before the hillshade is clipped into integer bounds, and is thus able to increase the dynamic range of the output image. We try multiple scaling factors, ensuring the mean value of the hillshade is near 128 (center of the distribution) and that the standard deviation is at minimum 45 (20% of the distribution). We begin with a best-guess zFactor, generate the hillshade, then update our guess based on how far the generated hillshade is from these target statistics. This re-guessing occurs up to 10 times, stopping at the first iteration where the new guess is the same as the last guess. The final zFactor output is then used in a second attempt at the matching process. If this dynamic scaling also fails to produce a successful match, we choose to fall back to a zFactor of $\frac{1}{2}$, suppressing smaller features which may have been overrepresented by a high-contrast scaling (see Section 4 for further discussion of hillshade improvement).

3.3. Matching Validation—MWIR LTM Onto VSWIR HVM³/M³

All four cases resulted in a correct match between the VSWIR reflected light images and simulated thermal images. For validating the time of matching, we test the ability to match modeled thermal images, produced at whole number hour intervals, to the VSWIR image. The modeled thermal image successfully matches at the closest hour angle of acquisition to the time of VSWIR acquisition. In most cases, the images also still match for a few hours outside the actual acquisition time of the VSWIR image. In one case, the match was only successful at the captured hour, which is likely due to relatively sparse features in the image, limiting key-point and match pair generation (Figure 8). The case with added noise behaved in the same manner as its equivalent no-noise simulated image in the matching. In all cases, this confirms the IMPPAIL procedure matches thermal and VSWIR captured at the same time, satisfying expectations for performance for simultaneous acquisitions during mission operations.

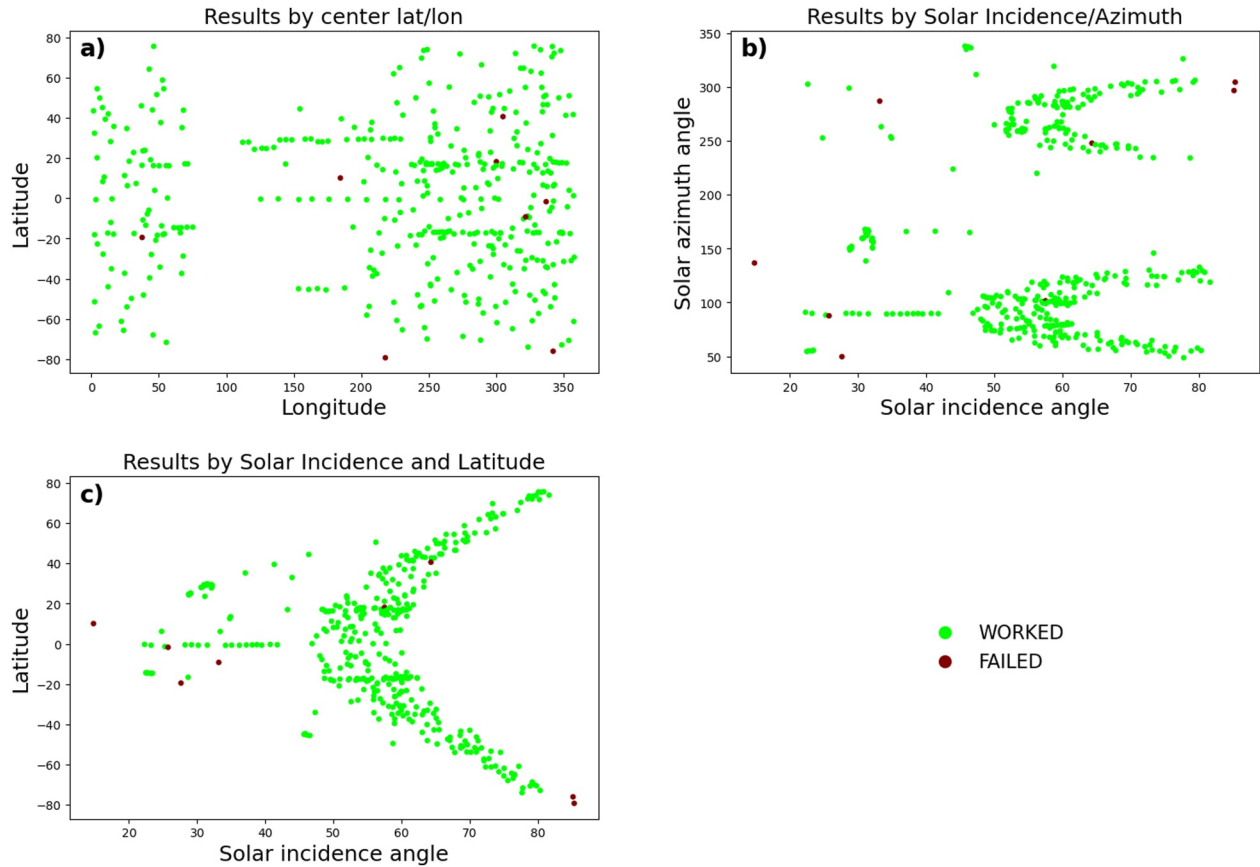


Figure 7. Success rate plotted by observation geometry for (a) longitude/latitude, (b) solar incidence/azimuth angles and (c) solar incidence/latitude.

3.4. Spatial Precision in Matching

Inspecting all successful matches of VSWIR and hillshade, we observe qualitatively near-perfect registration for each successfully matched image. We also quantify the matching accuracy by computing the mean pixel offset (using the L^2 norm) for each of the match pairs. We use the computed homography to transform the key-points from the M^3 image to the hillshade image and compute the mean offset of each pixel used in matching. Averaging over all match pairs within a given image pair, we find the range for mean accuracy on our data set is between 0.80 and 1.74 pixels. Across all test cases, the mean distance is 1.36 pixels. Given we are operating at a M^3 resolution




Figure 8. Thermal matching results. Green region represents when simulations successfully matched to the acquired M^3 image, while the cross marks represent the solar time when the M^3 image was actually acquired. A successful match at hour angle X is shown by a green bar including $X:00 \pm 30$ min.

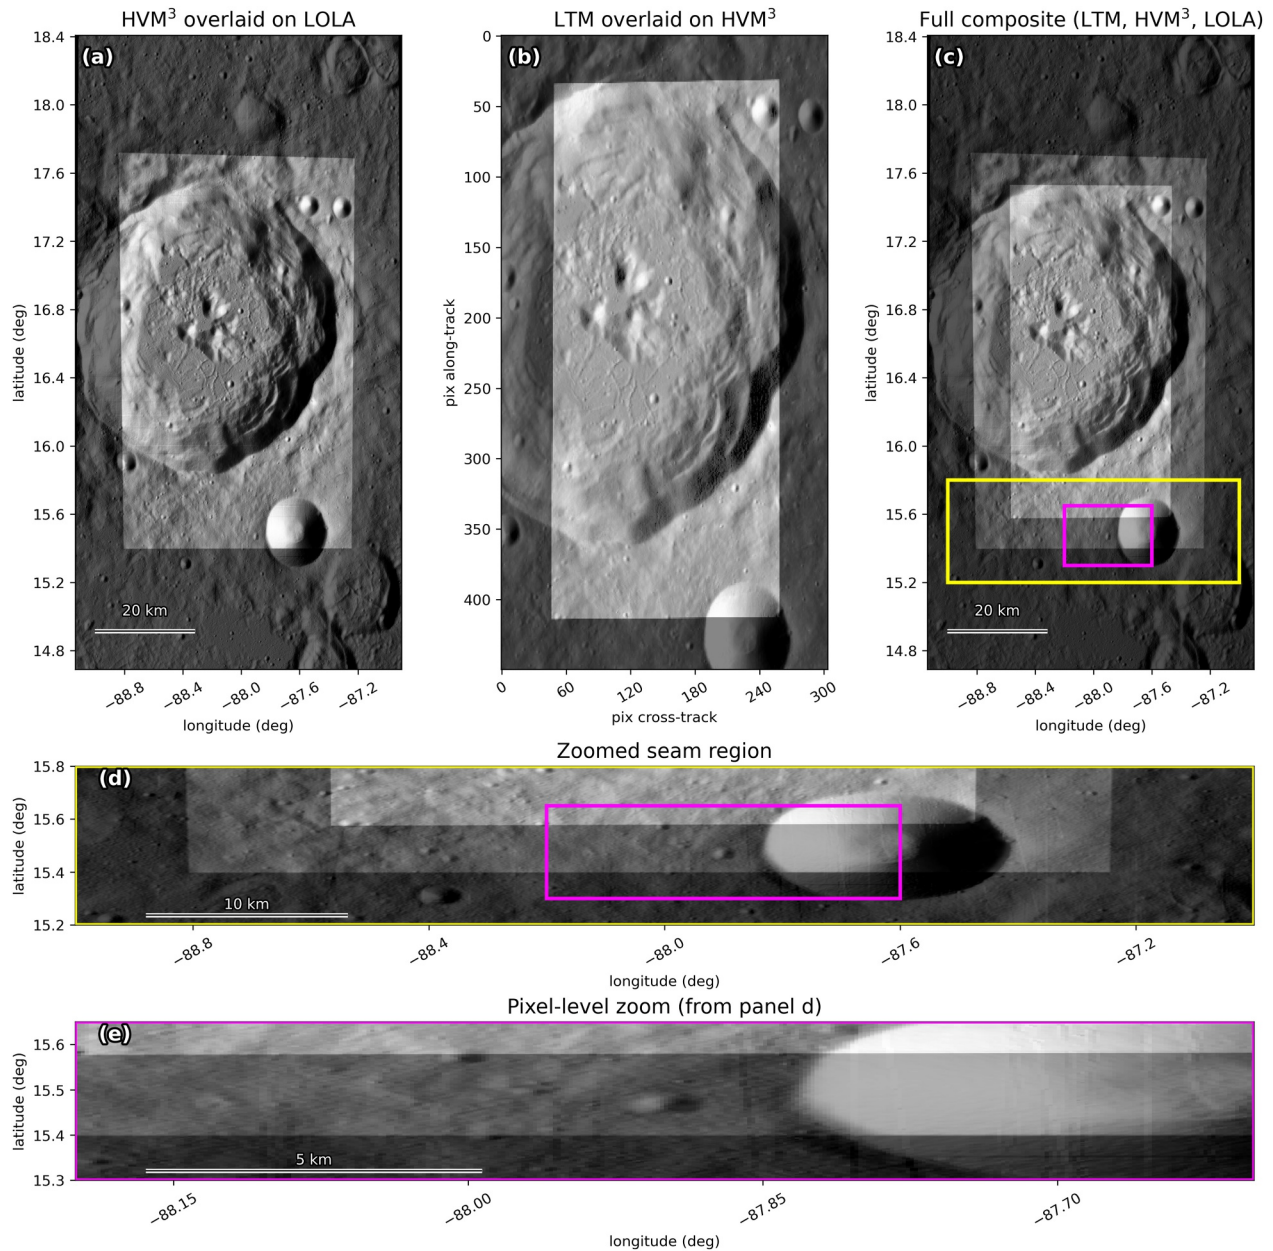


Figure 9. Final data products, displayed as blended overlay images. For the two intermediate products, VSWIR HVM³-analog data from M³ image ID M3G20090212T082712 is matched to (a) the LOLA hillshade and (b) the LTM simulated data is registered to HVM³. (c) Using the HVM³/LOLA registration, the registered LTM/HVM³ pair is overlaid onto the LOLA hillshade, resulting in a combined product including LTM, HVM³ and LOLA. The center image does not have associated latitude/longitude coordinates, as it is a match between the two images in pixel space. The matching backplane is used in the rightmost image, confirming accuracy in localization. (d) magnifies the seam at the top of the composite, allowing visual inspection of matching, with (e) further zooming to the worst-case seam region to show pixel-scale registration accuracy along the seam.

of 140 m/pix and a hillshade resolution of 60 m/pix, the same surface area will occupy 3 times as much pixel area in the hillshade, so a feature could feasibly match to the correct feature within $140/60 = 2.33$ pixels of the true feature. This gives a theoretical match radius of approximately 1.16 for true matches, close to our mean radius of 1.36 pix. This accuracy is demonstrated visually (Figure 9), where the “seam” formed by overlaying the three co-registered images matches with no edge aberrations.

3.5. Sensitivity Analyses

To this point, we have worked with the baseline uncertainties of the Lunar Trailblazer mission. We also check the stability of the algorithm under uncertainty above the nominal case. To simulate this, we grow the hillshade area about the center point, creating a larger region in which to localize the M^3 image. We use an area scaling where the area occupied by the detected image is halved, so the area of the hillshade relative to image is doubled. Maintaining the aspect ratio, we use a 91.1×159.4 km region (compared to 64.4×112.7 km) for the hillshade. Notably, of the eight misaligned images, seven of them are now marked successful, confirming that these cases with greater uncertainties can be successfully matched onto an appropriately large hillshade region (Section 3.1). The other one does not match because the M^3 localization still lies partially outside our match region, even once it is expanded by a factor of 2 (see Section 3.2). Thus, more positional uncertainty does not impact success, maintaining 396/404 successful matches even with expanded uncertainty, producing a consistent 98% success rate.

4. Discussion

Overall, the IMPPAIL results show the approach is ready for use on Lunar Trailblazer and other missions. Algorithm optimization and performance tests on data sets closely analogous to those which will be acquired during flight should enable similar accuracies for the mission. With 98% accuracy on VSWIR image matching in IMPPAIL, we have built upon the performance of earlier algorithms (Boardman et al., 2011). We show high-fidelity results on high-resolution images within a large uncertainty region, enabling use for a wide variety of missions to enable automated pixel-scale co-registration.

Throughout the development of IMPPAIL, we have prioritized the flexibility of all software written, ensuring that, once images are available, the software architecture will be ready to localize and process matches for Lunar Trailblazer. There has also been significant focus on making the pipeline work independently of input image type, allowing the processing approach to be used for many missions to come. This flexibility ensures that, while IMPPAIL software was built for Lunar Trailblazer, it can be applied for many future missions, enabling software-based solutions for hardware-based uncertainties in image placement. In proving registration entirely post-acquisition, we enable future missions to use similar architectures.

Future improvements to IMPPAIL must address the failure modes occurring at very low or high solar incidence angles. In our results, we achieved an 82% success rate (23/28) for images captured with an incidence angle below 40° or above 80° . Given the overall 396/404 match rate, this mode alone accounts for over half of the failures in our data set, even though it only accounts for $\sim 7\%$ of the data. This failure mode is likely due to the shaded relief algorithm details that lead to shadows becoming ill-defined at extreme angles, washing out potential features used in the matching process. In implementation, we expect raytraced hillshades (Gill, 2012) could somewhat mitigate this problem, where the rays will create shadows as they would be seen by the spacecraft, matching more closely to the image collected. However, in most cases, the GDAL hillshade procedure used in this paper can accurately capture the illumination of the scene and is easier to automate with higher speed and accessibility.

This version of IMPPAIL serves as a baseline estimate of success rate with Lunar Trailblazer-like data. The next step to validating these models is to perform the registration as is required by the mission, once the orbiter is actively collecting and acquiring data. In future use, we must also consider the multiple projection geometries of the mission. Currently, we use a global sinusoidal projection for all test cases, which works very well for low and mid latitude imagery ($< \sim 80^\circ$ latitude), but creates issues with polar geometries, where the matching should be switched to a polar stereographic projection or be locally map-projected about the center of the image. Ultimately, IMPPAIL will work on either data projection, producing the needed latitude-longitude backplanes for future processing steps. IMPPAIL is sufficiently generalizable to be adapted to any set of mission image data.

5. Conclusions

The iterative matching pipeline for post-acquisition image localization (IMPPAIL) has been successfully implemented and validated for use on the Lunar Trailblazer mission and with other image-based mission data sets. We demonstrated that accurate post-acquisition image geolocation can be successfully executed, without high-precision pointing and camera models, by using an algorithmic approach to identify match-pairs and transform

images to their correct latitude/longitude placement when a pre-existing planetary data set is available. The primary goal of mapping VSWIR and MWIR images to each other then to global topography maps was successfully performed with M³ VSWIR images, simulated thermal images, and LOLA/Kaguya digital elevation model data with uncertainty of less than 2 pixels in 98% of test images. We have also shown this success rate can be obtained with even higher initial placement uncertainties up to twice the nominal uncertainty, providing confidence even in more uncertain observing conditions. The supplementary IMPPAIL code files (Gauld et al., 2026a, GitHub/Zenodo) make this process usable and open, allowing for improved post-processing and colocalization for image data for other investigators working with different combinations of data sets.

Conflict of Interest

The authors declare no conflicts of interest relevant to this study.

Availability Statement

All data from M³ (Green et al., 2011; Lundeen, 2011) and the LOLA/Kaguya DEM (Barker et al., 2021) are publicly available through NASA's Planetary Data System. A code supplement detailing the algorithms applied is available at (Gauld et al., 2026a, GitHub/Zenodo), with data available at (Gauld et al., 2026b, CaltechDATA), which includes our simulated thermal data sets.

Acknowledgments

We are grateful to the Lunar Trailblazer Science Team for feedback and review throughout the development process of this software. We would also like to thank the NASA Planetary Data System for making data easily accessible for batch processing, and the developers of the OpenCV and OSGeo/GDAL packages for their implementations of various functions within the codebase. We also thank the reviewers whose comments improved this manuscript. Lunar Trailblazer was funded at Caltech campus under NASA contract 80MSFC19C0042. A portion of this research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0019F0079). The University of Oxford portion of this work was funded by the UK Space Agency.

References

- 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. <https://doi.org/10.1016/j.pss.2020.105119>
- Barker, M. K., Mazarico, E., Neumann, G. A., Zuber, M. T., Haruyama, J., & Smith, D. E. (2016). A new lunar digital elevation model from the lunar orbiter laser altimeter and SELENE terrain camera. *Icarus*, 273, 346–355. <https://doi.org/10.1016/j.icarus.2015.07.039>
- Bender, H. A., Smith, C. D., Ehlmann, B. L., Thompson, D. R., Vinckier, Q. P., & Mouroulis, P. (2022). Optical design and performance of the lunar trailblazer high-resolution volatiles and minerals moon mapper (HVM3). *Proceedings of SPIE*, 3. <https://doi.org/10.1117/12.2632552>
- Boardman, J. W., Pieters, C. M., Green, R. O., Lundeen, S. R., Varanasi, P., Nettles, J., et al. (2011). Measuring moonlight: An overview of the spatial properties, lunar coverage, selenolocation, and related level 1B products of the moon mineralogy mapper. *Journal of Geophysical Research*, 116(E00G14), E00G14. <https://doi.org/10.1029/2010JE003730>
- Bowles, N. E., Ehlmann, B. L., Evans, R., Warren, T. J., Eshbaugh, H. H., King, G., et al. (2026). The lunar trailblazer lunar thermal mapper instrument. *Journal of Geophysical Research*. <https://doi.org/10.1029/2025JE009333>
- Clark, R. N. (2009). Detection of adsorbed water and hydroxyl on the moon. *Science*, 326(5952), 562–564. <https://doi.org/10.1126/science.1178105>
- Colaprete, A., Schultz, P., Heldmann, J., Wooden, D., Shirley, M., Ennico, K., et al. (2010). Detection of water in the LCROSS ejecta plume. *Science*, 330(6003), 463–468. <https://doi.org/10.1126/science.1186986>
- Dickson, J. L., Ehlmann, B. L., Kerber, L., & Fassett, C. I. (2024). The global context camera (CTX) mosaic of Mars: A product of information-preserving image data processing. *Earth and Space Science*, 11(7), e2024EA003555. <https://doi.org/10.1029/2024EA003555>
- Edwards, C. S., Nowicki, K. J., Christensen, P. R., Hill, J., Gorelick, N., & Murray, K. (2011). Mosaicking of global planetary image datasets: 1. Techniques and data processing for thermal emission imaging System (THEMIS) multi-spectral data. *Journal of Geophysical Research*, 116(E10), E10008. <https://doi.org/10.1029/2010JE003755>
- Ehlmann, B. L., Bellerose, J., Lantoiné, G., Furlan, E., Scire, E., Fajardo-Acosta, S., et al. (2026). Lunar trailblazer spacecraft tracking and Mission recovery attempt: Characterization of status and behavior of a non-cooperative object in cis-lunar space. *Earth and Space Science*. <https://doi.org/10.1029/2025EA004734>
- Ehlmann, B. L., Klima, R. L., Seybold, C. C., Klesh, A. T., Bennett, C. L., Bowles, N., et al. (2026). The lunar trailblazer mission: Science motivation and implementation of a pioneering small satellite for lunar water and lunar geology in the NASA SIMPLEX program. *Journal of Geophysical Research*. <https://doi.org/10.1029/2025JE009300>
- Ehlmann, B. L., Klima, R. L., Seybold, C. C., Klesh, A. T., Au, M. H., Bender, H. A., et al. (2022). NASA's lunar trailblazer mission: A pioneering small satellite for lunar water and lunar geology. In *2022 IEEE aerospace conference (AERO), big sky* (pp. 1–14). <https://doi.org/10.1109/AERO53065.2022.9843663>
- Feldman, W. C., Maurice, S., Binder, A. B., Barraclough, B. L., Elphic, R. C., & Lawrence, D. J. (1998). Fluxes of fast and epithermal neutrons from lunar prospector: Evidence for water ice at the lunar Poles. *Science*, 281(5382), 1496–1500. <https://doi.org/10.1126/science.281.5382.1496>
- Fischler, M. A., & Bolles, R. C. (1981). Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6), 381–395. <https://doi.org/10.1145/358669.358692>
- Gauld, K. D., Dickson, J. L., Warren, T. J., Yang, J., Bowles, N., & Ehlmann, B. L. (2026a). Image localization: Code for the iterative matching pipeline for post-acquisition image localization (IMPPAIL). *LTB-Code/Image-Localization*, GitHub/Zenodo. <https://doi.org/10.5281/zenodo.18636928>
- Gauld, K. D., Dickson, J. L., Warren, T. J., Yang, J., Bowles, N., & Ehlmann, B. L. (2026b). Supporting data for post-acquisition image-based localization with application to the lunar trailblazer mission [Dataset]. *CaltechDATA*. <https://doi.org/10.22002/ad4b0-9cp74>
- GDAL/OGR Contributors. (2023). GDAL/OGR geospatial data abstraction software library. *Open Source Geospatial Foundation*. Retrieved from <https://gdal.org>
- Gill, K. M. (2012). Casting shadows: Shading digital elevation models using ray tracing. *Rivier University Academic Journal*, 8(1), 1–15.

- Green, R. O., Pieters, C. M., Mouroulis, P., Eastwood, M. L., Boardman, J. W., Glavich, T., et al. (2011). The moon mineralogy mapper (M³) imaging spectrometer for lunar science: Instrument description, calibration, on-orbit measurements, science data calibration and on-orbit validation. *Journal of Geophysical Research*, *116*(E10), E00G19. <https://doi.org/10.1029/2011JE003797>
- Hayne, P. O., Bandfield, J. L., Siegler, M. A., Vasavada, A. R., Ghent, R. R., Williams, J.-P., et al. (2017). Global regolith thermophysical properties of the moon from the diviner lunar radiometer experiment. *Journal of Geophysical Research: Planets*, *122*(12), 2371–2400. <https://doi.org/10.1002/2017JE005387>
- King, O., Warren, T., Bowles, N., Sefton-Nash, E., Fisackerly, R., & Trautner, R. (2020). The Oxford 3D thermophysical model with application to PROSPECT/luna 27 study landing sites. *Planetary and Space Science*, *182*(2020), 104790. <https://doi.org/10.1016/j.pss.2019.104790>
- Li, S., Lucey, P. G., Milliken, R. E., Hayne, P. O., Fisher, E., Williams, J. P., et al. (2018). Direct evidence of surface exposed water ice in the lunar polar regions. *Proceedings of the National Academy of Sciences*, *115*(36), 8907–8912. <https://doi.org/10.1073/pnas.1802345115>
- Li, S., & Milliken, R. E. (2017). Water on the surface of the Moon as seen by the moon mineralogy mapper: Distribution, abundance, and origins. *Science Advances*, *3*(9), e1701471. <https://doi.org/10.1126/sciadv.1701471>
- Lowe, D. G. (1999). Object recognition from local scale-invariant features. *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, *2*, 1150–1157. <https://doi.org/10.1109/ICCV.1999.790410>
- Lowe, D. G. (2004). Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, *60*(2), 91–110. <https://doi.org/10.1023/B:VISI.0000029664.99615.94>
- Lundeen, S. R. (2011). *Moon mineralogy mapper calibration data record*. NASA Planetary Data System. CHI-ORB-L-M3-4-L1B-RADIANCE-V3.0, CH1M3_0003. <https://doi.org/10.17189/1520248>
- McEwen, A. S., Eliason, E. M., Bergstrom, J. W., Bridges, N. T., Hansen, C. J., Delamere, W. A., et al. (2007). Mars reconnaissance orbiter's high resolution imaging science experiment (HiRISE). *Journal of Geophysical Research*, *112*(E5), E05S02. <https://doi.org/10.1029/2005JE002605>
- Paige, D. A., Foote, M. C., Greenhagen, B. T., Schofield, J. T., Calcutt, S., Vasavada, A. R., et al. (2010). The lunar reconnaissance orbiter diviner lunar radiometer experiment. *Space Science Reviews*, *150*(1–4), 125–160. <https://doi.org/10.1007/s11214-009-9529-2>
- Pieters, C. M., Goswami, J. N., Clark, R. N., Annadurai, M., Boardman, J., Buratti, B., et al. (2009). Character and spatial distribution of OH/H₂O on the surface of the moon seen by M³ on Chandrayaan-1. *Science*, *326*(5952), 568–572. <https://doi.org/10.1126/science.1178658>
- Robinson, M. S., Brylow, S. M., Tschimmel, M., Humm, D., Lawrence, S. J., Thomas, P. C., et al. (2010). Lunar reconnaissance orbiter camera (LROC) instrument overview. *Space Science Reviews*, *150*(1–4), 81–124. <https://doi.org/10.1007/s11214-010-9634-2>
- Rubanenko, L., Venkatraman, J., & Paige, D. A. (2019). Thick ice deposits in shallow simple craters on the moon and mercury. *Nature Geoscience*, *12*(8), 597–601. <https://doi.org/10.1038/s41561-019-0405-8>
- Smith, D. E., Zuber, M. T., Neumann, G. A., Lemoine, F. G., Mazarico, E., Torrence, M. H., et al. (2010). Initial observations from the lunar orbiter laser Altimeter (LOLA). *Geophysical Research Letters*, *37*(18), L18204. <https://doi.org/10.1029/2010GL043751>
- Spudis, P. D., Bussey, D. B. J., Baloga, S. M., Cahill, J. T. S., Glaze, L. S., Patterson, G. W., et al. (2013). Evidence for water ice on the moon: Results for anomalous polar craters from the LRO Mini-RF imaging radar. *Journal of Geophysical Research: Planets*, *118*(10), 2016–2029. <https://doi.org/10.1002/jgre.20156>
- Sunshine, J. M., Farnham, T. L., Feaga, L. M., Groussin, O., Merlin, F., Milliken, R. E., et al. (2009). Temporal and spatial variability of lunar hydration as observed by the deep impact spacecraft. *Science*, *326*(5952), 565–568. <https://doi.org/10.1126/science.1179788>
- Thompson, D. R., Ehlmann, B. L., Green, R. O., Allen, G. D., Bender, H., Copley-Woods, D., et al. (2026). Calibration and performance of the high-resolution volatiles and minerals moon mapper (HVM3). *Earth and Space Science*. <https://doi.org/10.1029/2025EA004456>
- Zuber, M. T., Smith, D. E., Zellar, R. S., Neumann, G. A., Sun, X., Katz, R. B., et al. (2010). The lunar reconnaissance orbiter laser ranging investigation. *Space Science Reviews*, *150*(1–4), 63–80. <https://doi.org/10.1007/s11214-009-9511-z>