Automatic Mapping of Small Lunar Impact Craters Using LRO-NAC Images

J. H. Fairweather[1] , **A. Lagain[1]** , **K. Servis[1,2]** , **G. K. Benedix[1,3,4]** , **S. S. Kumar[5]** , and **P. A. Bland[1,3]**

[1]Space Science and Technology Centre, School of Earth and Planetary Sciences, Curtin University, Perth, WA, Australia, [2]CSIRO, Pawsey Supercomputing Centre, Kensington, WA, Australia, [3]Department of Earth and Planetary Sciences, Western Australia Museum, Welshpool, WA, Australia, [4]Planetary Science Institute, Tucson, AZ, USA, [5]Department of Geology and Geophysics, Indian Institute of Technology, Kharagpur, India

Abstract Impact craters are the most common feature on the Moon's surface. Crater size–frequency distributions provide critical insight into the timing of geological events, surface erosion rates, and impact fluxes. The impact crater size–frequency follows a power law (meter-sized craters are a few orders of magnitude more numerous than kilometric ones), making it tedious to manually measure all the craters within an area to the smallest sizes. We can bridge this gap by using a machine learning algorithm. We adapted a Crater Detection Algorithm to work on the highest resolution lunar image data set (Lunar Reconnaissance Orbiter-Narrow-Angle Camera [NAC] images). We describe the retraining and application of the detection model to preprocessed NAC images and discussed the accuracy of the resulting crater detections. We evaluated the model by assessing the results across six NAC images, each covering a different lunar area at differing lighting conditions. We present the model's average true positive rate for small impact craters (down to 20 m in diameter) is 93%. The model does display a 15% overestimation in calculated crater diameters. The presented crater detection model shows acceptable performance on NAC images with incidence angles ranging between ~50° and ~70° and can be applied to many lunar sites independent to morphology.

Plain Language Summary The Moon's surface is covered in impact craters and recording their spatial density gives researchers the ability to study the geological evolution of our satellite. Analyzing craters helps in determining the physical properties of planetary surfaces and how/if impact rates change over time. These analyses rely on recording spatial densities for numerous surfaces, which has been achieved for craters >1–2 km on the Moon. Manually counting the smaller craters, which number in the hundreds of millions, is a daunting task. We adapted a Crater Detection Algorithm and applied it to the highest resolution lunar imagery data set. We describe our method for gathering, reformatting, and detecting craters across lunar images down to 20 m in diameter. The detection model performance was quantitatively evaluated across six different regions, each with different terrain and lighting conditions. Comparison between manually mapped craters and detections from our model allows us to conclude that the model has an acceptable performance in detecting fresh to moderately degraded craters of all sizes, down to 20 m in diameter, when compared to other studies. Automated crater detection complements manual counting methods and aids in unlocking secrets of the Moon's surface.

1. Introduction

Mapping and recording structural features across celestial bodies is one of the ways researchers gather spatial and temporal context about evolution of our solar system. As a readily available structure across a host of bodies, impact craters give us the interface into understanding and interpreting extraterrestrial geological events (see, Fassett, 2016; Hartmann, 1966; Ivanov et al., 2002; Neukum et al., 2001; Öpik, 1960; Shoemaker & Hackman, 1962; Young, 1940). Rocky bodies, such as the Moon, act as cosmic records for the solar system, where impact craters accumulate on the surface over time (Hartmann, 1965; Melosh, 1989). The development of modern space technologies and satellite systems has made locating and mapping impact craters with ever increasing detail a familiar task for many planetary scientists (Fassett, 2016; Ivanov et al., 2002). Craters of all sizes inform us of the physical characteristics of the impacted terrains such as their density, porosity, and composition (e.g., Melosh, 1989; van der Bogert et al., 2017), as well as areas of scientific interest (potential landing sites). For example, the Moon has varying amounts of regolith with different degrees of thicknesses, and by measuring impact craters, we can model and estimate those regolith thicknesses (Bart et al., 2011; Rajšić et al., 2021; Stopar

Validation: J. H. Fairweather
Writing – original draft: J. H. Fairweather
Writing – review & editing: A. Lagain, K. Servis, G. K. Benedix, S. S. Kumar, P. A. Bland

et al., 2012, 2017; Wilcox et al., 2005). Specifically, small impact craters (<1 km) aid in determining the rate of erosion (Fassett & Thomson, 2014; Soderblom, 1970) and calculating surface retention ages (Qiao et al., 2018).

Crater chronology methods, based on recording cratering densities across different surfaces, give the ability to assign model ages to different geological units and events (Baldwin, 1965; Hartmann, 1965; Moore et al., 1980; Neukum et al., 2001; Shoemaker & Hackman, 1962; Williams et al., 2018). Surfaces with more craters are older than less cratered surfaces, as they have been exposed longer (Neukum et al., 2001; Shoemaker & Hackman, 1962). This relative age relationship becomes a model age when defined by the radiometrically dated returned lunar samples (Apollo, Luna, and now Chang'e). Though this method is not without error and requires the ability to accurately see and map impact craters within an area (Baldwin, 1964; Povilaitis et al., 2018; Robbins et al., 2014). Secondary craters, which can dominate a crater population (<1 km), are formed from ejected material which fails to reach escape velocity. These craters typically occur in clusters radiating away from the primary impact (McEwen & Bierhaus, 2006). This ejected material contains crucial information about the primary impactor and impacted surface (Lagain, Benedix, et al., 2021). Ideally when calculating model ages, secondary craters should not be used, though it can be difficult to differentiate secondaries from small primary impacts (McEwen & Bierhaus, 2006; Povilaitis et al., 2018; Stadermann et al., 2018). Small craters across the Moon's surfaces are in a state of equilibrium (<1 km for post-Nectarian surfaces), meaning the number of craters being obliterated is proportional to the crater density of the impacted surface of a given age (Hartmann, 1971; Minton et al., 2019; Moore et al., 1980; Povilaitis et al., 2018; Schultz et al., 1977; Xiao & Werner, 2015). Creating a data set of small impact craters gives the opportunity to gather more detail into the equilibrium and retention states across younger terrains (Minton et al., 2019). Recording the small crater population in and around larger craters and landforms also gives spatial context to the secondary clusters and transported material (McEwen & Bierhaus, 2006; Schultz et al., 1977).

All such crater analyses rely on a complete cratering record (all craters to a given size counted) across a range of diameters (Crater Analysis Techniques Working Group, 1979; Xiao & Strom, 2012). It is well established that the cumulative number of impact craters on a planetary surface follows an inverse power law relationship (Hartmann, 1965; Ivanov et al., 2002; Robbins et al., 2018; Shoemaker & Hackman, 1962). This relationship means that to count ever-smaller craters across a surface reliably, one must count exponentially increasing numbers of them. The task of counting craters has historically (pre-2000) been restricted to manual mapping methods, first using maps, rulers, and pencils (see, Baldwin, 1964, 1987a, 1987b; Hartmann, 1965, 1977), and now using Geographical Information System (GIS) tools (see, Kneissl et al., 2011 [for the GIS Tool]; Head et al., 2010; Povilaitis et al., 2018; Robbins, 2019 [for the Moon]; Lagain, Bouley, et al., 2021; Robbins & Hynek, 2012 [for Mars]; Herrick et al., 2018 [for Mercury]; Zeilnhofer, 2020 [for Ceres]; Liu et al., 2018 [for Vesta]). These manual methods have resulted in hundreds of thousands to millions of craters counted on different bodies. However, studies of human attention span have shown that the ability to identify any feature consistently and reliably significantly decreases with time spent on the task (Langner & Eickhoff, 2013; See et al., 1995). This, and personal interpretation, leads to variability of impact crater identification between researchers (Robbins et al., 2014). This can potentially introduce errors to total crater numbers, crater sizes, and degradation measurements. In many cases, the errors can be accounted for.

Robbins (2019) compiled a comprehensive, manual, global database of lunar impact craters (~1.3 million entries), complete down to 1–2 km in diameter using the Lunar Reconnaissance Orbiter-Wide Angle Camera (LRO-WAC) imagery (resolution = 100 m/px) and Digital Elevation Models (DEMs) from Lunar Orbiter Laser Altimeter (LOLA)/SELENE Kaguya (resolution = 59 m/px at ±60° latitude and 100 m/px across ±90° latitude; Robinson et al., 2010; Smith et al., 2010). This vast data set made clear the level of manual mapping that can be accomplished using the latest satellite data sets and GIS. At present, the highest resolution images available for the Moon are the Lunar Reconnaissance Orbiter-Narrow-Angle Camera (LRO-NAC) images (NAC images, at 0.25–2 m/px), which allows scientists to map and record hundreds of thousands of decimeter-sized craters within a single NAC image (~5 km by ~25 km; Robinson et al., 2010). As the orbiter is still actively imaging the lunar surface and the image data set is continually updated, we have the potential to completely map the lunar surface. However, creating a global record of small (decimeter diameter range) craters on the Moon using manual mapping techniques is inconvenient (Cadogan, 2020; Hashimoto & Mori, 2019; Lee & Hogan, 2021). Automated crater mapping techniques can provide the bridge into the global analysis of decimeter-sized craters.

There has been great progress in automated lunar crater detection within the last decade, allowing researchers to increase the scale and consistency in which we map craters. Comprehensive reviews by both DeLatte et al. (2019) and Stepinski et al. (2012) show that the field of automating crater detections has been investigated since the early 1980s. The advancement of modern Convolutional Neural Networks (CNNs), a type of deep machine learning algorithm that focuses on object detection in visual data sets (see, Goodfellow et al., 2016, Chapter 9 for a review), has provided leap forward for planetary science. The heavily cratered surfaces of Moon and Mars have been key targets for generating and testing automated crater data sets. Current published automated data sets (discussed further below) use both DEM-based and optical imagery-based methods to high degrees of accuracy (measured in true detection rate or Recall) when compared against manual crater data sets (e.g., Head et al., 2010; M. Jia et al., 2020; Povilaitis et al., 2018; Robbins, 2019). Hybrid automated methodologies (crater data sets based on DEM and optical imagery), such as Salamunićcar et al. (2012, 2014), had success in using an automatic Crater Detection Algorithm (CDA) for the lunar surface, producing a published catalog of 19,396 detections—which was incorporated into the Salamuniccar et al. (2012) LU60645GT catalog (see, LU78287GT catalog, 78,287 craters complete to ~8 km). Their method was one of the few to utilize both Hough transforms (a different type of algorithm that detects crater boundaries) and crater shape interpolation, which combines detections from a modified LOLA DEMs and WAC images to generate a crater data set. Their data set had a detection rate of 83% (Salamunićcar et al., 2012) and ~85% (Salamunićcar et al., 2014). DEM-based methods, such as Silburt et al. (2019), developed a CDA trained on the LOLA/Kaguya DEM. In their study, the team detected 15,658 craters between ~2 and ~90 km (across eastern side of the Moon, 60°E–180°E), with a high detection rate of 92%. Silburt et al. (2019)'s CDA stands out as their team showed the transferability of their lunar detection model to different celestial bodies, such as Mercury (using MESSENGER DEM, 665 m/px) with good preliminary results (also see, Lee [2019] for the model's modification for Mars craters). Wang et al. (2021) also developed a DEM-based CDA for use across LOLA DEM, which detected craters with three-dimensional morphological characteristics (such as rim height, interior slope, and depth). Their crater data set (LU1319373) consists of a massive 1 million craters >1 km, with a detection rate of ~85%. Strictly image-based methods, such as Hashimoto and Mori (2019), aimed at detecting small craters (<50 pixels) across the lunar south pole using LRO-NAC imagery. Their two-staged method cuts the image data into grids/tiles, runs a semantic segmentation algorithm (a CNN algorithm that outputs a modified image instead of a list of detections) followed by an object detection algorithm to gather the position of the craters within in the grid, with a detection rate 80.5%. Another image-based CDA by Cadogan (2020), tested automatic detection of small craters (>2.5 m) by using six overlapping LRO-NAC images. His analyses were conducted over the Apollo, Luna, Lunokhod, and Ranger sites, with the most extensive analysis over the Lunokhod 1 site, where >300,000 craters over 2.5 m in diameter were automatically detected with a calculated detection rate 86%. Many these studies recommend using a DEM (Di et al., 2014; Silburt et al., 2019; Wang et al., 2021) or a hybrid-based method (Salamunićcar et al., 2011; Yang et al., 2020) for making a global crater data set. While DEMs have a good consistent global coverage they lack the resolution needed to map craters to the subkilometer scale (Urbach & Stepinski, 2009). The benefit of using an image-based CNN object detection algorithm for the Moon (as opposed to DEM-based CNN) is the availability of LRO-NAC high-resolution image data. The global coverage of the image data, compared to that of the lower resolution global DEMs, allows for far greater effectiveness in small crater (<1 km) detection.

In our approach, we aim for a reliable detection model for impact craters, making a data set complete to the small (20 m) crater scale, wherever suitable NAC images are available, using an object detection, image-based, machine learning algorithm. We accomplish this by adapting an existing object detection CNN model designed to globally identify impact structures on high-resolution planetary images of Mars (see, Benedix et al., 2020; Lagain, Benedix, et al., 2021; Lagain, Servis, et al., 2021) to the Moon. A trained CDA has the potential to address the large variabilities in consistency, time-consumption, and reproducibility, which influences manual crater counts. This paper presents the methodology and evaluation of our automated CDA across LRO-NAC images.

2. Methodology

The machine learning framework we adapted in this present study, originally outlined in Benedix et al. (2020), works in three phases (Figure 1). The first phase (Figure 1a) preprocesses the input NAC imagery by changing image projection and file format. The second phase (Figure 1b) is a series of supervised learning and validation cycles, producing a trained detection model. Finally, the third phase (Figure 1c) uses that detection model
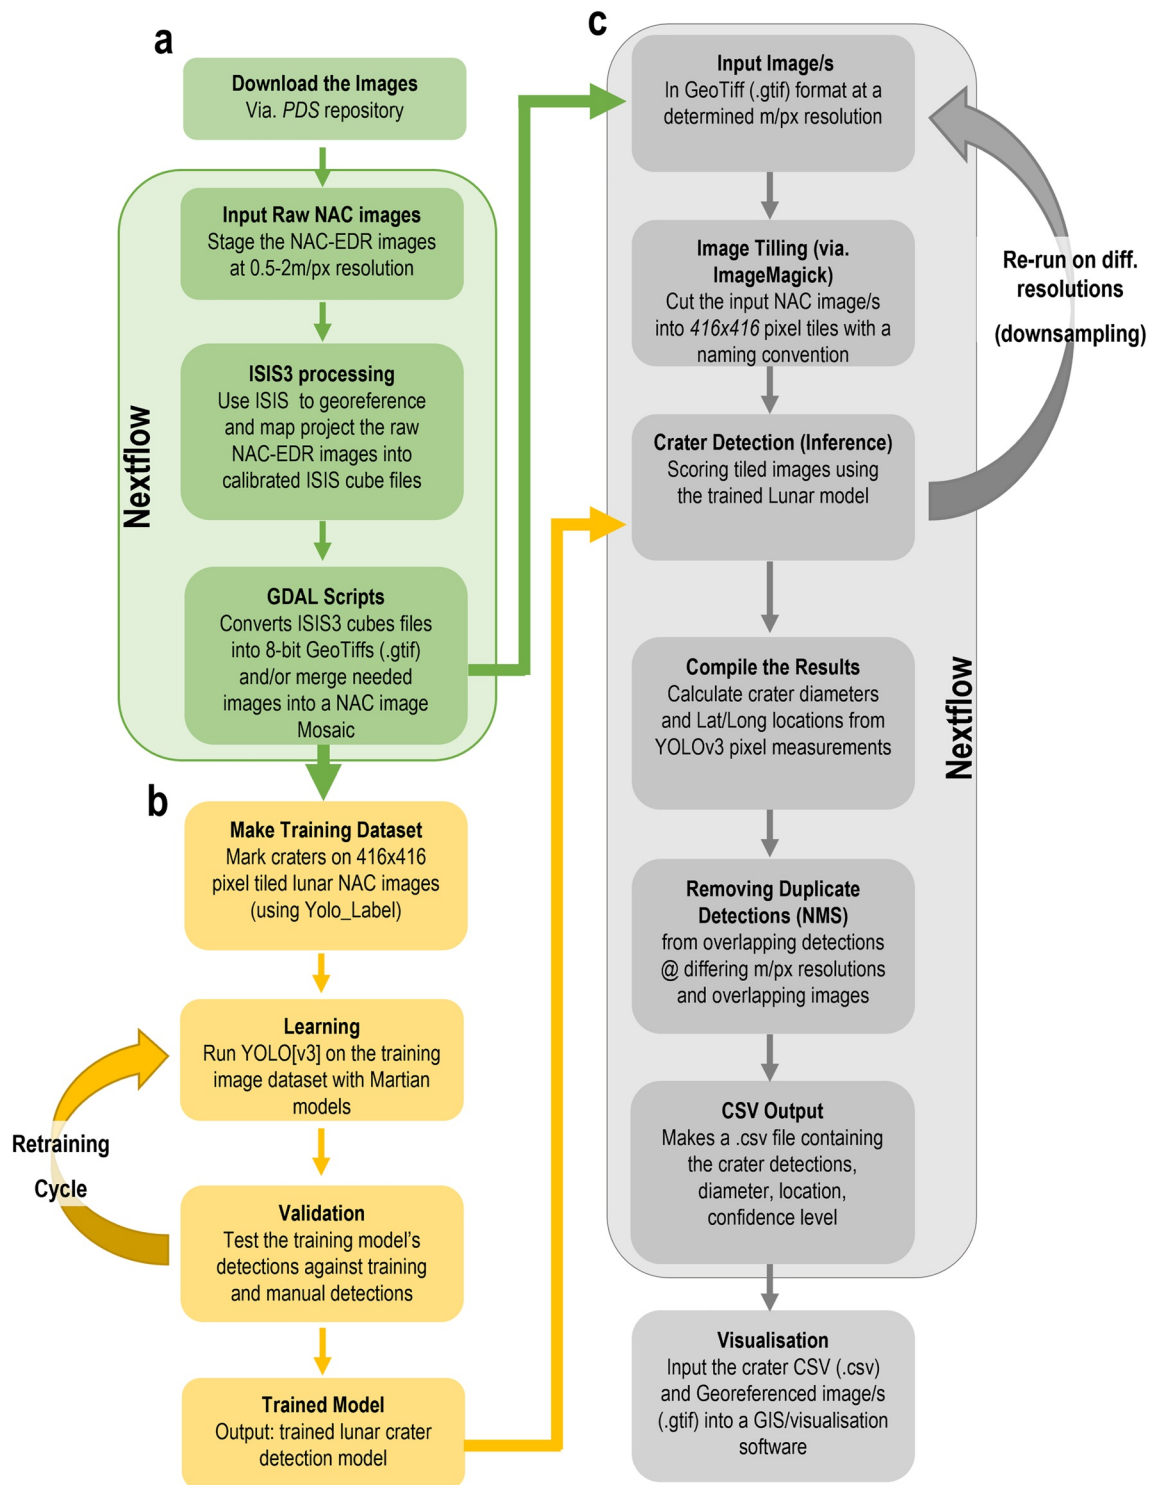
Figure 1. Overview flow diagram depicting the processing pipeline for our Crater Detection Algorithm (CDA). (a) The preprocessing needed to georeference the raw Narrow-Angle Camera (NAC) images ready for algorithm training and inference; (b) the workflow for model training; (c) the processes for impact crater detection across NAC images.

to identify craters across a suite of input images to produce a list of detections. The subsequent detections are analyzed for duplicates, and where found are removed through a Non-Maximum Suppression (NMS) approach using detection location and confidence value (see, Benedix et al., 2020; Figure 1c, and further explained in

Section 2.3.1). In this study, we utilized an open-source software workflow framework called Nextflow (Di Tommaso et al., 2017) to concatenate and run phase a and c (Figure 1). The Nextflow program aided in orchestrating the processes for each of the data-intensive pipelines of this CDA.

2.1. Phase 1: Image Preprocessing

The image data set used for crater detection, the LRO-NAC data set, is the highest resolution global image repository (0.25–2 m/px spatial resolution) available for the Moon; however, it is not available in the georeferenced format (e.g., geotiff/gtif) required to determine the coordinates of detected craters (see, Benedix et al., 2020; Lagain, Servis, et al., 2021). The LRO-NAC images were downloaded from the publicly available Planetary Data System (https://pds.nasa.gov/). Relevant images were map projected using United States Geological Survey (USGS) Integrated Software for Imagers and Spectrometers (ISIS) and converted into the GeoTiffs using GDAL (see Supplemental Text 1 in Supporting Information S1 for details). A supercomputer cluster at Pawsey Supercomputing, Perth, Western Australia (https://pawsey.org.au/) is utilized to process each image. Using this resource, a single NAC image (∼5 km by ∼25 km, ∼250 MB) takes only seconds to process; magnitudes faster compared to an average personal computer.

2.1.1. Image Selection Rationale

The performance of any object detection algorithm depends significantly on the quality of the training data set and the images analyzed (DeLatte et al., 2019). The more representative the selection of training images, the more robust the CDA results will be. Therefore, it is imperative that training occurs on NAC images with favorable lighting conditions (i.e., illumination angles that do not hinder crater recognition) and across different lunar terrains. Terrain selection is quite straight forward. The surface of the Moon is divided into two broad terrains: Highlands and Mare, where each terrain has a different impact crater morphology (Shoemaker, 1964; Wilhelms, 1987). The lunar Highlands are older, anorthositic in composition, with a rugged mountainous morphology (Wilhelms, 1987); the Mare terrains are younger, smooth flood basalt plains primarily found within huge impact basins on the lunar nearside (Hiesinger et al., 2011; Wilhelms, 1987).

Illumination angle selection took a little more work. The NAC data set (from https://wms.lroc.asu.edu/lroc/search) is very extensive, consisting of 3,000,000+ images (each ∼200–500 MB in size) with more routinely added each year, resulting in different areas imaged many times over with different angles of sun lighting. Since launch (2009), the LRO spacecraft has taken a wealth of images over the same regions with differing conditions (refer to Figure S1 in Supporting Information S1), giving us the ability to choose the images that best illuminate the different terrains for crater detection. The image lighting conditions are defined by the solar incidence angle (Figure 2), with vertical illumination (noon) at 0° and horizontal illumination (sunrise/sunset) at 90°. Low incidence angles (<40°) have high contrasts, where fresh, bright crater ejecta can mask smaller or mature craters (Figure 2a). Similarly, lunar topography interpretation is affected by higher incidence angles (>80°), due to large shadows cast by peaks/rims covering swaths of the lunar surface (Figure 2c). Therefore, images with incidence angles between ∼50° and 70° are the most reliable for identifying all features important for counting craters (Cadogan, 2020; Head et al., 2010; Robbins, 2019). However, Richardson et al. (2022) analyzed crater populations acquired on different incidence angle images and found that fewer craters are identified on images with incidence angles lower than 58°, therefore, we focus training our network using NAC images that have incidence angles >∼58° and evaluate the performance with respect to the lighting conditions.

2.2. Phase 2: Network Training

This CDA is a CNN image object detection algorithm, using an Ultralytics implementation of You Only Look Once version 3 (YOLOv3-ultralytics) as the base neural network architecture (Jocher, 2021; Redmon & Farhadi, 2018; Redmon et al., 2016). We trained the computer through a process that required a series of supervised learn and validate cycles—where each learn cycle compares the results against a control data set (DeLatte et al., 2019). In our case, the control is a ground truth data set of manually marked craters across images. To train efficiently, we chose and processed NAC images that best represented the lunar surface and the variability in craters.
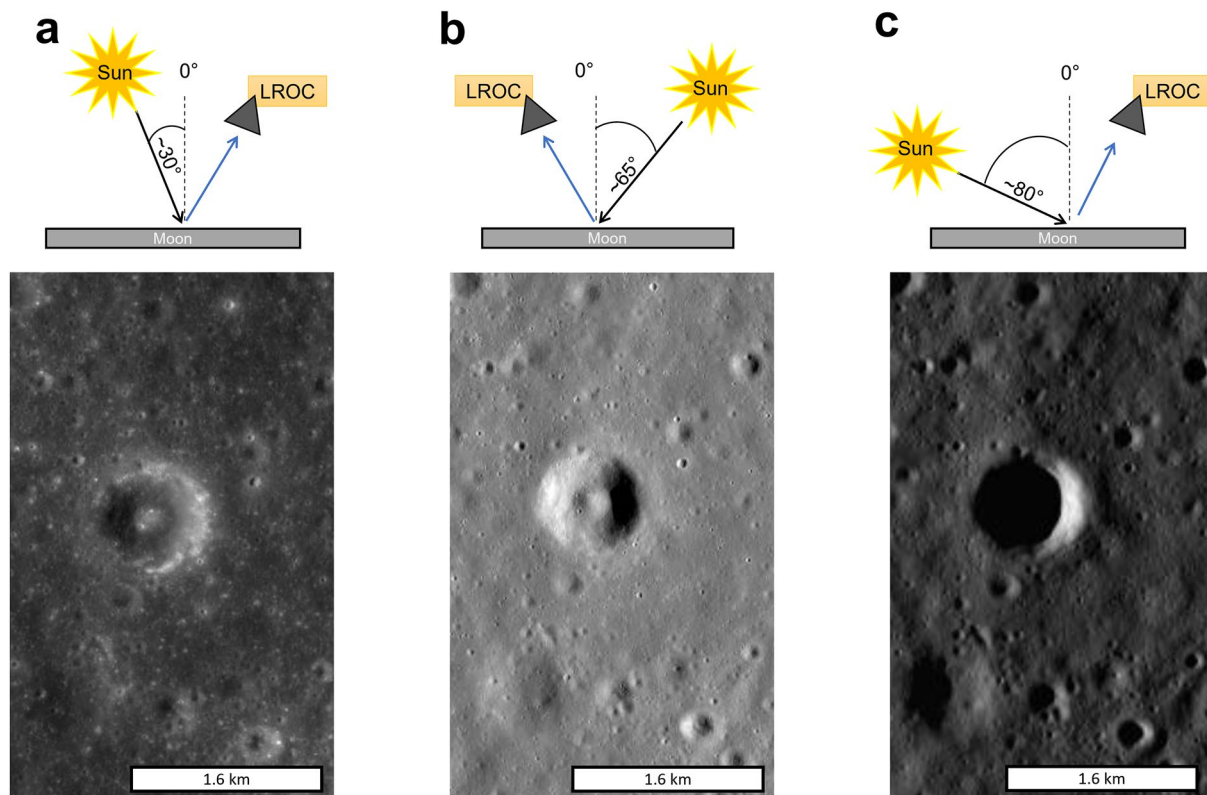
Figure 2. Example of Narrow-Angle Camera (NAC) images illustrating the effect of three different lighting conditions on the features (lighting condition is indicated by the images incidence angle in degrees). (a) NAC image M1282458049L, incidence angle of 27.7°; (b) NAC image M1108275380R, incidence angle of 66.5°; (c) NAC image M1277755323L, incidence angle of 83.2°. The 800 m centroid crater in the three images is located at 0.68°S 0.91°E.

2.2.1. The Training NAC Images

We chose a suite of NAC images for training based on their variability in crater form and lighting conditions. The Apollo 14 landing region (Figure 3) hosts a sufficient variability of mappable craters across both highland and mare terrains (specifically, Fra Mauro highlands, Mare Insularum, and Mare Cognitum). Highland terrains are older than mare terrains, therefore hosting more craters, where the complex topography affects the crater shape (Robbins et al., 2014). This difference in shape means there is more visual variability between highland craters; thus, more marked craters are required to capture that variability. As the Apollo 14 landing site region is an area of interest, there is larger pool of NAC images and mosaics to choose from. Moreover, when detecting craters across the Moon down to <1 km, the morphological differences between craters located on the near and on far sides are negligible (Daubar et al., 2014).

The chosen image data set used to train the detection model were composed of 16 NAC images. Of these, 14 were from areas within a ~300 km radius of the Apollo 14 landing site. The 14 images, grouped by terrain, are (see Figure 3) *Mare Insularum* (M1127206936L, M119964604LE, M119998529L, M1320200353L, and M1363614226L); *Mare Cognitum* (M1215555724L); and the *Fra Mauro Highlands* (M102265088L/R and M1096608496L/R, M1249655817L, M1310797424L, M131772598L, and M1363656518L). The two remaining NAC images sampled *Oceanus Procellarum* (M1305167327L) and *Terra Vitae* (M1378458043L; Figure 3). These two additional areas ensure that the CDA has a broader detection scope, and the subsequent detection model is more representative of the range of terrains of the Moon. These test images have incidence angles ranging between 45.67° and 81.81°. Ten of the images have the narrower range of 64.76°–76.23°, a range shown to be favorable for detecting craters (Cadogan, 2020; Richardson et al., 2022; Schultz et al., 1977).

2.2.2. Creating the Training Data Set

For the algorithm to efficiently learn and validate on the NAC images, the images were tiled (cut the image into 416 × 416 pixel thumbnails) using ImageMagick editing package (Still, 2006). This resulted in hundreds of
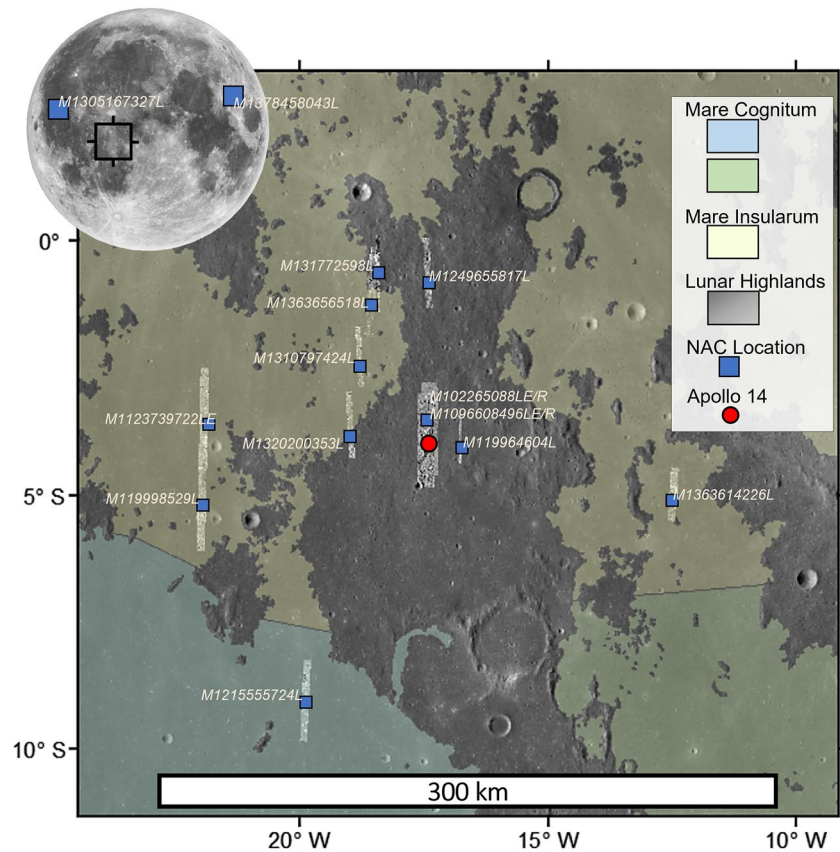
Figure 3. Lunar terrains (Highland [*Fra Mauro highlands*] and Mare [*Mare Insularum* and *Mare Cognitum*]) surrounding the Apollo 14 landing site (red marker, 3.65°S 17.47°W), and the image footprint locations and ID #'s of the 16 Narrow-Angle Camera (NAC) images used for training (blue marker and text). Image locations not within the Apollo 14 region are denoted in the inset globe in the upper left. Base image is the Wide Angle Camera (WAC) global mosaic (Speyerer et al., 2011) and mare boundaries are from Nelson et al. (2014).

image tiles per NAC image. As the pixel dimensions of the tiles are fixed and that NAC resolution can vary from 0.25 to 2 m/px, the maximum sized crater we could mark would be 416 pixels diameter, that is, between 104 and 832 m depending on the image resolution (if the crater was situated perfectly in the middle of the tile). To extend the range of sizes detected, the pixel resolutions of the 16 NACs were downsampled to 2, 10, and 20 m/px (maximum crater diameter of 832, 4,160, and 8,320 m within a tile, respectively). Marking was carried out on semirandomly chosen image tiles. This resulted in 248 tiles from the 16 NAC images where we marked impact craters manually. We used Yolo_Label software tool to label the tiles with boxes. These bounding boxes fit the crater rim as accurately as possible to get a precise crater size measurement (refer to Figure S2 in Supporting Information S1 for an example of a marked tile). Because object detection algorithms can only reliably detect objects larger than ~10 pixels (Wang et al., 2020), we marked impact craters down to 10 pixels wide on the training tiles. Note that this size threshold constitutes a lower limit from which impact craters are accurately recognized and measured manually (Robbins et al., 2014). The craters marked within the tiles are mostly simple craters, as the simple-complex transition diameter on the Moon is 15 km (Krüger et al., 2018). Thus, the CDA is optimized to detect simple craters. After following through with this process, our resulting ground truth training data set was composed of 43,402 impact craters (~200 craters per tile on average) with diameter range of ~5 m to ~3 km (all tiles used available in supplemental Data Set S3). This quantity of marked craters is significantly higher relative to other ground truth data sets for similar published algorithms: Hashimoto and Mori (2019) with 4,967; Benedix et al. (2020) with 1,762; Yang et al. (2020) with 14,406 craters; and Lagain, Benedix, et al. (2021) with 2,142 craters.
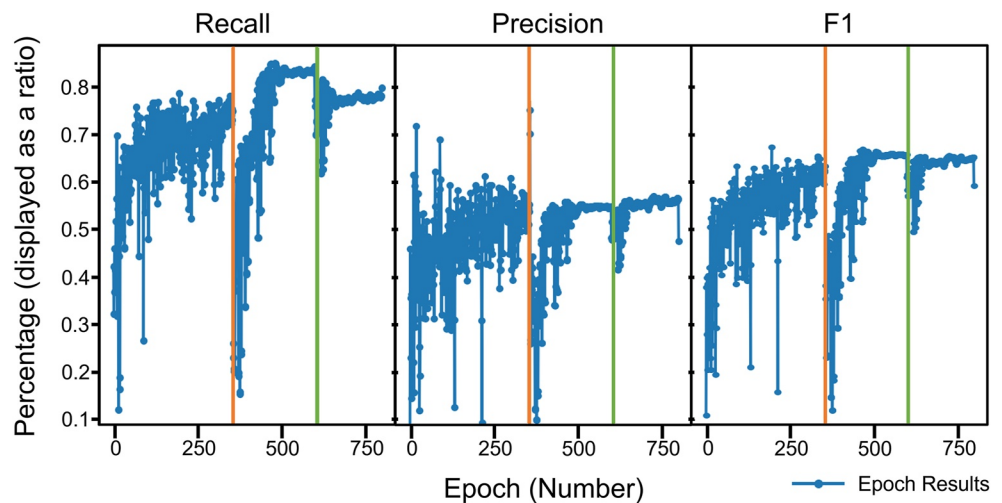
Figure 4. Training metrics for the lunar crater detection model in YOLOv3. The orange line marks the start of the first training session (300–600), from the model initially trained on High-Resolution Imaging Science Experiment (HiRISE) images (0–300, from Benedix et al., 2020; Lagain, Benedix, et al., 2021; Lagain, Servis, et al., 2021), the green line marks the start of the second training session (600–800). As the training and validation cycles increase, the values of Recall, Precision, and F1 converge to roughly constant values. See the main text for the definition of the recall, precision, and F1 score percentage values.

2.2.3. Training the Network

Training in YOLOv3-ultralytics entailed running a series of learning and validation cycles, which requires marked objects on images for learning (learning data set) and objects on images to compare against (validation data set). The outcome after a specified number of cycles is a detection model. To make the learning and validation data sets, we followed the approach outlined in Benedix et al. (2020), where whole ground truth data set (248 marked tiles) was randomly subdivided into the learning and validation data sets with a ratio of 4:1. The training process is a series of finite cycles called epochs (Bhavsar & Ganatra, 2012; Goodfellow et al., 2016; Zhang, 2010). Each epoch comprises a learning stage followed by a validation stage. Consecutive epochs adjust image augmentations, to slightly change the visual characteristics (i.e., rotation, translation, shear, and scale) on the image. This step removes the significance of sunlight direction (solar azimuth) across the data set but is beneficial as it effectively extends the training set to more crater variations. There is a possibility of undertraining (i.e., the model is very generalized and detects objects that generally look like craters) or overtraining (i.e., the model detects very specific looking craters). The key is finding a balance (DeLatte et al., 2019; Silburt et al., 2019) to get a model with homogeneous performance.

2.2.4. Transfer Learning

Optimizing the lunar detection model required two training sessions, using a transfer learning approach. The first session ran from 300 to 600 epochs (Figure 4) and trained on 188 NAC image tiles with a narrow range of incidence angles (56°–73°). This particular training session used a Martian crater detection model (see, Lagain, Benedix, et al., 2021; Lagain, Servis, et al., 2021), that used High-Resolution Imaging Science Experiment (HiRISE, 25 cm/px) images as the benchmark starting point. YOLO needs a starting point, so providing a model already optimized for impact craters helps the training algorithm look for circular crater-like features immediately, greatly increasing the efficiency of initial learn-validate cycles. The second training session was on 248 image tiles, which contained the original 188 tiles plus an additional 60 tiles and had a broader range of incidence angles (45°–81°). This session used the first model as a starting point and continued to an additional 200 epochs (a total of 500 epochs for the Moon, Figure 4).

Three important metrics derived from the results of the training sessions are used to evaluate performance of the algorithm; Recall, Precision, and F1 score (metrics also used by Lee, 2019; Lee & Hogan, 2021; Salamunićcar et al., 2014; Silburt et al., 2019). These are calculated after each epoch and training session (Figure 4 and see Figure S3 in Supporting Information S1 for all metrics). Each metric is defined based on a confusion matrix

(see, Figure S4 in Supporting Information S1; True Positive (TP), True Negative (TN), False Positive (FP), False Negative (FN), and Ground Truth data (GT)) and defined as follows:

1. Recall is the ratio of True Positive to Ground Truth values (TP/GT), which indicates the proportion of manual detections the algorithm correctly identified (i.e., True Positive detection rate);
2. Precision (TP/(TP + FP)), indicates the proportion of true craters the CDA detected (i.e., how many of the total positive detections are True Positive detections); and
3. F1 score (TP/(TP + 1/2(FP + FN))), which is the harmonic mean (numerical average) between Precision and Recall.

High False Positive and False Negative values result in lower Precision and Recall values, respectively. The F1 score serves as an indication to the overall error (lower values indicate more error). These performance indicators are discussed in detail in Section 3.1. Precision, Recall, and F1 values also indicate if the model is overtrained or undertrained. Low Recall and Precision would indicate the model is undertrained, as it fails to recognize craters. Whereas an overtrained model performs exceptionally well on the training data and poor across an evaluation (test) data set (images the model did not train on). The final choice of which model to use is based on the highest performance results; in this case, epoch 790 has the best balance between these metrics.

2.3. Phase 3: Running the Trained CDA on NAC Images

Inference, the process of running input data (preprocessed NAC images) through a trained model to produce a scored output, is the final phase of the CDA. As the volume of input data can reach a few terabytes if entire NAC images are analyzed, the utilization of YOLO as the CNN architecture is particularly relevant for our purpose because it subdivides the input images. Each is scored only once using weighted bounding boxes around identified craters, thus making the detection process extremely fast compared to other CNN architectures (Redmon & Farhadi, 2018). Each detection receives a confidence value, calculated by the CNN, which is a ratio comparing the detected crater to the trained crater model (Redmon & Farhadi, 2018). Only detections with a confidence value higher than 0.3 are kept. The resulting detection data, in the form of bounding boxes per image tile, are reassembled back into a complete NAC image. The detections, however, are still in a YOLO format, with the box location defined by pixel percent ratio. For example, in an image tile that is 100×100 pixels, a random point on that image would be [x] 0.6 (60% or 60 px), [y] 0.3 (30% or 30 px). The CDA has code to convert the YOLO dimensions into lunar longitude–latitude coordinates using the embedded geographical data of the original NAC image.

As mentioned in Section 2.2.2, when detecting craters across a NAC image tile (416×416 pixels), the largest possible complete crater the algorithm can mark is limited by the resolution. To mediate this, we used an approach widely used in object detection (e.g., Benedix et al., 2020; Lee, 2019; Salamuniccar et al., 2014), where we iteratively downsampled each NAC image to a lower resolution for consecutive executions of the CDA. This pyramidal approach allows the CDA to detect craters across a large dynamical range of sizes (Figure 5). To build crater detections across a range of diameters, the CDA evaluated three different versions of the same NAC image: one at 2 m/px (the lowest homogeneous resolution that can be applied to the global NAC data set; Figure 5a); the next at 10 m/px (Figure 5b); and finally at 20 m/px (Figure 5c). Since the smallest size for an accurate CDA object is 10 pixels (Wang et al., 2020), our minimum crater sizes are 20, 100, and 200 m, respectively. The downsampled pixel resolutions allow different crater diameter targets, meaning the algorithm runs across an image multiple times (similar to the method outlined in Cadogan [2020]), collating the crater detections across a range of diameters. This process is parallelized to run across multiple NAC images which greatly speeds up the processing. The computation time for inference depends on the size of the image and computer setup, but generally analysis of a single NAC image covering a ~125 km^2 area is ~30 s on the supercomputer cluster at Pawsey Supercomputing Centre.

2.3.1. Detection Data Postprocessing

Running the CDA on downsampled and/or overlapping images results potentially in the multiple detections of the same craters. It is important to remove the duplicate detections (see Figure S5 in Supporting Information S1) but their size and location vary slightly due to differences in conditions between NAC images. To remove duplicates, we use the same method as Benedix et al. (2020): a spatial statistical-based approach called NMS, that uses
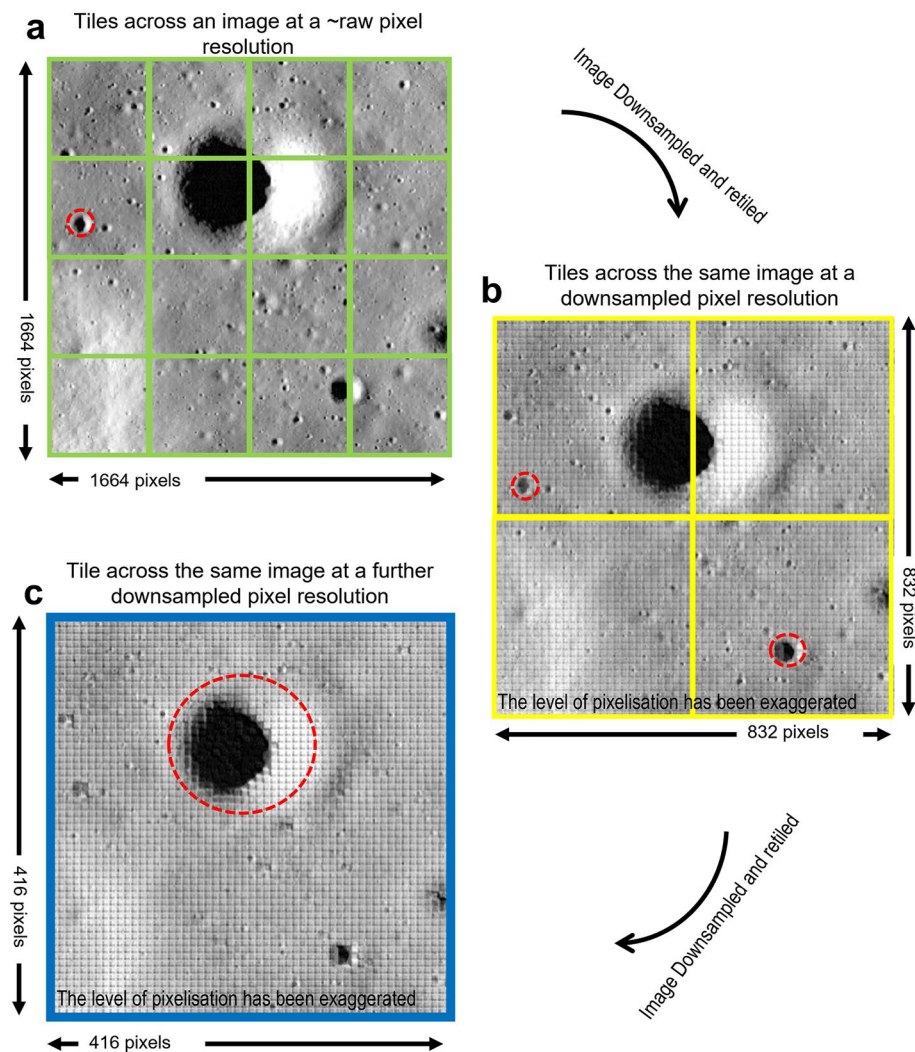
Figure 5. Effect of downsampling and then tiling a Narrow-Angle Camera (NAC) image for impact crater detection (base image is a square clip of NAC image M1096608496L, Fra Mauro highlands). (a) The green, 16-box-grid represents 416 × 416 pixels tile boundaries generated at the image's ~raw pixel resolution, note that at this resolution there are many significant craters that are bisected by tiling; (b) the yellow 4-box-grid represents the tiles generated at a downsampled pixel resolution; (c) the final large blue-box represents tile boundaries generated at a further downsampled pixel resolution, note that no craters are cut by tile lines, but there is trade-off in image detail. The level pixelization in (a) and (b) has been exaggerated for illustrative purposes. These illustrate the three pixel-resolution levels an input image would go through for crater detection. The red-dashed circles in each image outline the craters that can be completely detected by the Crater Detection Algorithm (CDA) at that resolution. Note, when downsampling and then tiling an image you allow bigger structures to fit within a single tile boundary (c) and greatly decrease the chance of features being truncated.

Intersection over Union (IoU) values to determine duplicates within a data set of spatial information (Figure 6). This technique removes all overlapping detections, keeping only the one with the highest confidence value. Specifically, the IoU value is a ratio describing how much one object's area overlaps with another (Figure 6b). The higher the ratio value (e.g., 0.9), the more of a match they are and therefore are most likely duplicates (Figure 6c). We used an IoU value of 0.3 as the threshold where a crater might have been multiply detected. This process compared each detected crater's area against all other detected craters nearby and determines IoU values for those overlapping relationships. This process repeated for every crater until it compared all craters. Theoretically, this has potential to remove nonduplicates, if the overlap is higher than the threshold. Though, the condition where two real craters of similar enough size, to have an overlap over the IoU threshold (>0.3), is negligible when compared to real duplicates influencing the data.
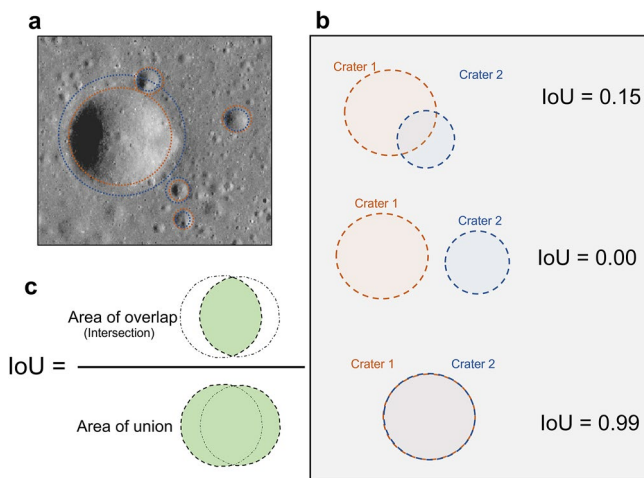
Figure 6. Intersection over Union (IoU) in crater duplicate removal. (a) Example of Crater Detection Algorithm (CDA) duplicates in the region of overlap between Narrow-Angle Camera (NAC) image M1157749492L and M1123590516L; (b) the IoU equation; and (c) visualization of the overlap relationship between two craters.

3. Evaluation of the Crater Detection Algorithm

3.1. Training Sessions

Across both training sessions, the final Recall, Precision, and F1 values were high, with an instability during the first ~100 epochs of each session (Figure 4). The instability reflects the algorithm learning, with a general upward trend to higher values as the model learns. The Recalls were notably high with end value of ~0.86 in session 1 and ~0.81 in session 2. This ~0.05 change is a function of the additional marked craters across a wider incidence angle range (45°–81°). The enhanced training set significantly increases the CDAs ability to (a) utilize more of the NAC data set and (b) detect more craters across a more extensive scope of lunar surfaces (global coverage). We consider this slight loss of performance acceptable with respect to the increasing amount of data where our model can return accurate results.

On the other hand, the Precision values are consistent across both sessions sitting at ~0.55, with a slight increase in the second session to ~0.56 (Figure 4). In the case of crater detection on ultra-high-resolution images, Precision can be a misleading, but informative, metric. The relatively low values are not always indicative of poor performance, as these values reflect the limitations in the completeness of ground truth training data set and image resolutions. As described earlier, our 43,402-crater ground truth data set is not complete for very small craters (<10 px/5–10 m). Therefore, during validation, the CDA finds it has detected a (very small) crater that has not been manually marked. The resulting detection is flagged as a False Positive (detection of a false crater), lowering the precision. This trend is noted in other CDA studies such as Silburt et al. (2019), where their Precision values were in the mid-50s (53% on a validation data set, and 56% on a test data set) for similar reasons.

3.2. The Evaluation Data Set

To evaluate the performance and accuracy of the CDA, we executed the model on 12 regions of interest from six semirandomly chosen NAC images (Figure 7), where each image chosen represented one of the two major lunar terrains over a range of lighting conditions (41°–67°). There were two areas per NAC image, one specifically for comparing smaller craters (>20–100 m) and another for comparing larger craters (>100 m to 1 km; see Table 1 and Data Set S3 for images of each area). This evaluation consists of three analyses, which quantify how the detection model performs in detecting craters of increasing dimeter. (a) The quantification of Recall, Precision, and F1 score for all detections across different diameter ranges, (b) the computation of similar metrics based on crater degradation discrimination, and (c) the accuracy on the crater size estimation.

For coherence, we input three groups of lighting conditions: IA-1 ("good/favorable" lighting for crater detection, 66.9°, 66.7°), IA-2 ("medium" lighting, 56.4°, 51.7°), and IA-3 ("poor" lighting, 48.9°, 41.8°) (Figure 7). Running the CDA and NMS over these six images (~900 km^2) took ~30 min and detected ~1.8 million craters (~10 m to 3 km). Furthermore, we instigated a 1 km diameter cutoff to set a maximum limit for evaluation, as we do not wish to compare against larger craters (>1 km) which many prior algorithms have already accomplished (e.g., Silburt et al., 2019; Wang et al., 2021; Yang et al., 2020). This translates to a comparison of ground truth craters to CDA-detected craters between >20 m and 1 km in diameter. Two ground truth data sets each focusing on a different crater diameters allowed evaluation of the CDA. The first is a "Small Crater Data Set" that compares 2,780 CDA detections against 2,159 manually mapped craters. This data set consists of marked craters with diameters of 20–500 m (±10%) across six 1–9 km^2 areas. The second data set is a "Large Crater Data Set" consisting of 1,540 CDA-detected craters and 1,343 manually mapped craters. This data set focused on evaluating larger craters with diameters of 100 m to 1 km (±10%) across six 20–150 km^2 areas (see supplemental Data Sets S1–S3).
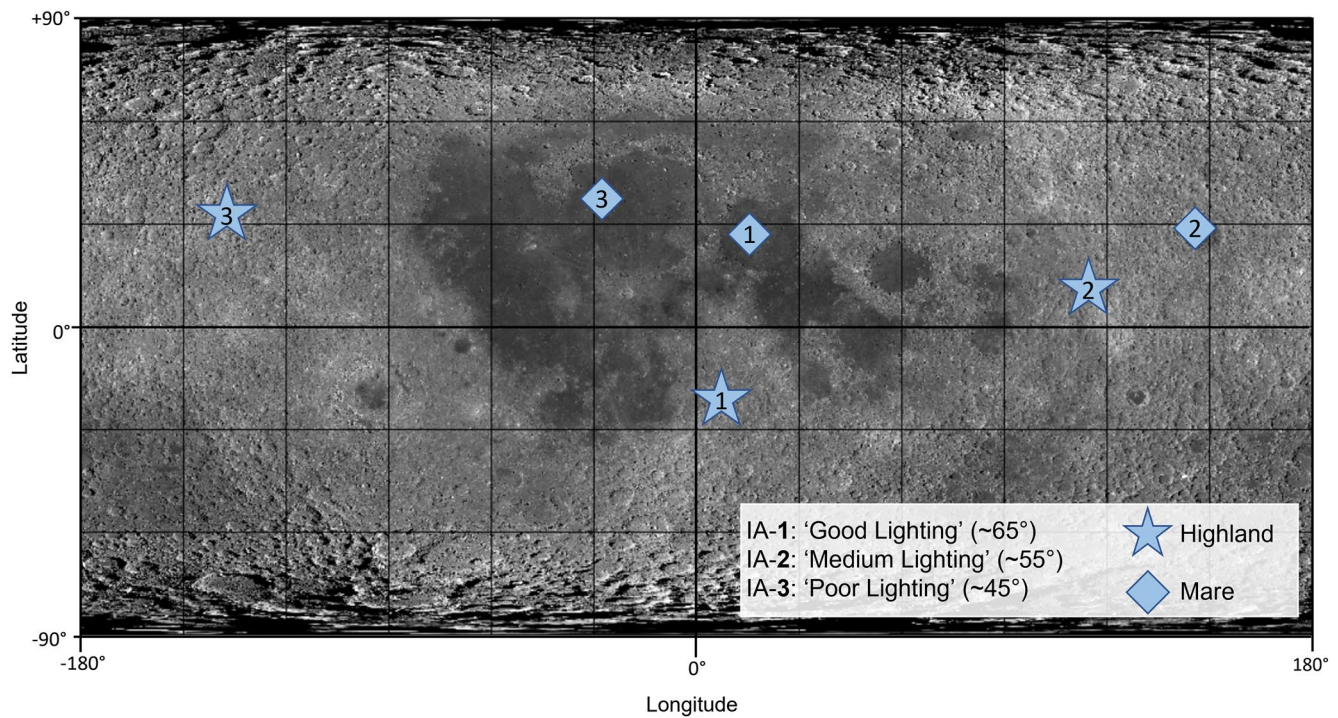
Figure 7. Locations of the Narrow-Angle Camera (NAC) images used in the evaluation (in a simple cylindrical projection). Blue stars are Highland terrains, blue diamonds are Mare terrains, the numbering (1–3) indicates the lighting condition group: IA-1 (66.9°,66.7°), IA-2 (56.4°, 51.7°), and IA-3 (48.9°, 41.8°).

3.3. CDA Performance

Recall, Precision, and F1 scores, for each evaluation test, were calculated based on the values from a confusion matrix presented in Figure S4 in Supporting Information S1. The Small Crater Data Set (20–500 m, Table 2) had an overall average Recall of 0.93, Precision of 0.66, and F1 score of 0.77. Across the images with favorable lighting conditions (IA-1), the CDA performed well, with an average recall of 0.90, precision of 0.71, and F1 score of 0.79. The CDA also performed similarly across the medium lighting condition (IA-2). However, the CDA performance decreased significantly on images with poorer lighting conditions (IA-3), specifically the F1 score, with an average of 0.72. This difference being due to the difficulty of consistently identifying craters in the high-contrast lighting conditions (noted by Cadogan, 2020; Robbins et al., 2014; Wilcox et al., 2005).

The Large Crater Data Set (100 m to 1 km, Table 2) had, comparatively, a lower overall average Recall of 0.89, similar Precision of 0.67, and a lower F1 score of 0.74. The values of the Recall, Precision, and F1 scores across the Large Crater Data Set was more varied, with high F1 values (0.84) across all craters (≥100 m) and lower values (0.62, 0.77) across the larger craters (≥300 m, ≥500 m). These errors stem from the CDA's crater diameter calculation and are exaggerated due to the larger diameters (further discussed in Section 3.3.2). Interestingly, across both crater data sets (Table 2), there are no significant differences in the numerical performance (Recall, Precision, and F1 score) between the near side or far side Highland and Mare terrains.

In summary, the average Recall, Precision, and F1 scores of our CDA across NAC images with favorable and medium lighting conditions (IA-1/IA-2/ Highland/Mare only) are 0.93, 0.70, and 0.79, respectively, for the Small Crater Data Set, and 0.86, 0.68, and 0.74 for the Large Crater Data Set. This is in line with the recall results of recently published lunar image-based CDAs such as Hashimoto and Mori (2019) with ~0.81 and Cadogan (2020) with ~0.86.

Table 1

Image Information of the Six NAC Images Used in Our Evaluation

Image ID	Depicted terrain	Incidence angle (°)	Centroid long/lat (decimal °)
M1338833866L	Near side Highland	66.9°	6.52°/−10.62°
M1325197569L	Far side Highland	56.4°	115.94°/9.99°
M1288076949L	Far side Highland	48.9°	−142.06°/30.39°
M1320016983L	Near side Mare	66.7°	7.48°/27.18°
M112963850L	Far side Mare	51.7°	151.64°/29.05°
M1361391010L	Near side Mare	41.8°	−28.12°/36.84°

Table 2
Summary of the Performance Metrics for the "Small Crater Data Set" and "Large Crater Data Set"

| | | IA-1 | | IA-2 | | IA-3 | | |
| | | 66.7° | 66.9° | 51.7° | 56.4° | 41.8° | 48.9° | |
Metric	Diameter (m)	Near side Mare	Near side Highland	Far side Mare	Far side Highland	Near side Mare	Far side Highland	Average
Small Crater Data Set								
Recall	≥20	0.98	0.97	0.88	0.90	0.86	0.92	0.92
	≥40	0.90	0.95	0.94	0.96	0.76	0.92	0.90
	≥60	0.84	0.92	1.00	1.00	0.91	1.00	0.94
	≥80	0.80	0.89	1.00	1.00	1.00	1.00	0.95
	≥100	0.83	0.89	1.00	1.00	1.00	1.00	0.95
Precision	≥20	0.76	0.68	0.84	0.90	0.70	0.79	0.78
	≥40	0.68	0.68	0.68	0.82	0.63	0.51	0.67
	≥60	0.77	0.54	0.50	0.64	0.53	0.50	0.58
	≥80	0.75	0.66	0.67	0.80	0.50	0.50	0.65
	≥100	0.76	0.76	0.40	0.67	0.63	0.60	0.64
F1	≥20	0.85	0.80	0.86	0.90	0.77	0.85	0.84
	≥40	0.78	0.79	0.79	0.88	0.69	0.66	0.77
	≥60	0.81	0.68	0.67	0.78	0.67	0.67	0.71
	≥80	0.77	0.76	0.80	0.89	0.67	0.67	0.76
	≥100	0.79	0.82	0.57	0.80	0.77	0.75	0.75
Large Crater Data Set								
Recall	≥100	0.91	0.85	0.97	0.95	0.91	0.90	0.91
	≥300	0.80	0.83	1.00	0.67	1.00	0.89	0.87
	≥500	0.80	0.85	-	0.78	1.00	1.00	0.88
Precision	≥100	0.74	0.87	0.75	0.80	0.70	0.78	0.77
	≥300	0.47	0.63	0.40	0.73	0.42	0.41	0.51
	≥500	0.29	0.92	-	0.88	1.00	0.60	0.74
F1	≥100	0.82	0.86	0.84	0.87	0.79	0.84	0.84
	≥300	0.59	0.72	0.57	0.70	0.59	0.57	0.62
	≥500	0.42	0.88	-	0.82	1.00	0.75	0.77

Note. The values are grouped based on their lighting condition, respective terrain, and diameter (the results for each cell are in Table S1 in Supporting Information S1).

3.3.1. Effects of Crater Degradation in CDA Performance

To evaluate the effectiveness of the CDA across different levels of degradation, we classified craters >100 m based on their degradation state. We assigned each crater a degradation class, A, B, or C, based on visual criteria and descriptions used by Pohn and Offield (1970), Trask (1971), Stopar et al. (2012), and Mahanti et al. (2018). The description of each class is (Figure 8) A: craters with sharp rims, steep slopes, crisp internal shadows, and visible rocky ejecta; B: craters with smoothed rims, shallow slopes, internal shadow, and little ejecta; and C: craters with ill-defined rims, very shallows slopes, and little internal shadows (Figure 8). We did not include craters degraded beyond these classifications. False Positives are impossible to calculate, as we are only comparing to a specific attribute of the ground truth crater data set throughout this analysis. The aim was to quantify the detection rate (Recall) of the CDA in detecting actual craters across three classes of degradation. We can only determine if the CDA failed or succeeded in detecting marked degraded craters.
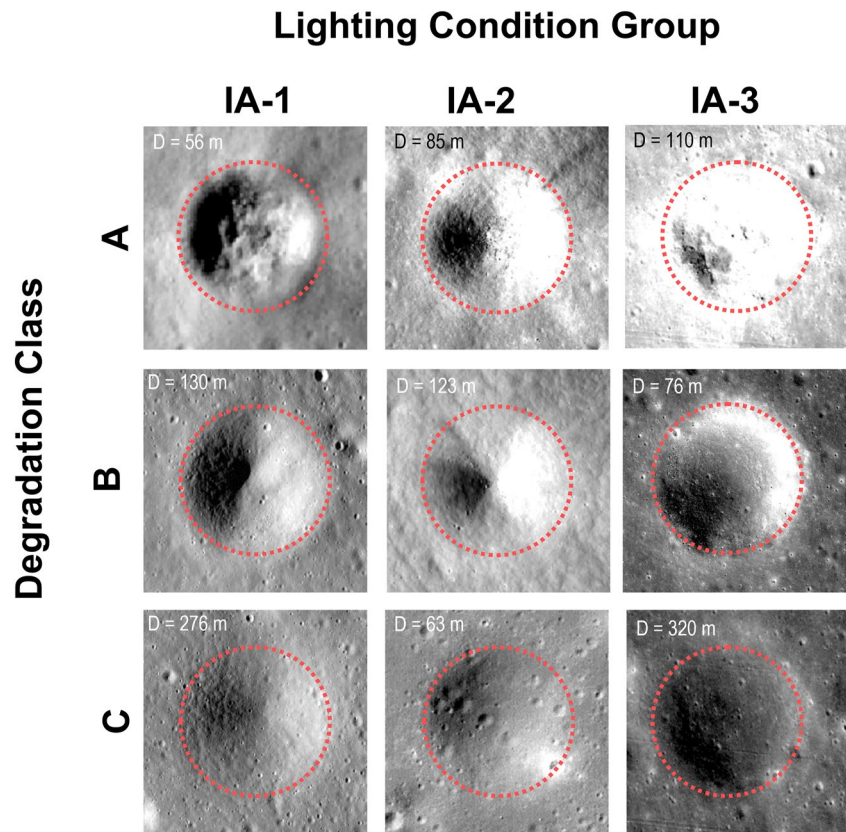
Lighting Condition Group

Figure 8. Examples of the crater degradation classes (A/B/C) for each lighting condition (IA-1/IA-2/IA-3). Red-dashed circles denote the crater boundaries. Each crater image is taken directly from the Narrow-Angle Camera (NAC) images used within this study.

Results show that our model is excellent at detecting class-A and class-B craters (Figure 8 and Table 2). The average Recall across these two classes sits at 0.99 (Table 3). There are also no significant differences between the Recall values across the lighting conditions (IA-1/2/3) and terrains. However, the Recall notably decreased (−0.15) across highly degraded (class-C) craters (Figure 8) with an average Recall of 0.85 (Table 3). There is more variability in recording the size for very degraded craters, for both humans (Fassett & Thomson, 2014; Robbins et al., 2014) and machines (Delatte et al., 2019; Y. Jia et al., 2020; Wang et al., 2020). For the CDA, the detection variability of degraded craters ranges from 0% to 30% (Table 3), which is roughly within the range of manual data sets (see, Robbins et al., 2014).

3.3.2. Crater Diameter Error Estimation

We evaluated how effective the CDA is at quantitatively determining the diameter of the detected craters. This analysis included over a total 3,261 true positive crater detections (2,058 from the Small Crater Data Set, and 1,203 from the Large Crater Data Set), comparing crater diameters from the ground truth (D_{GT}) with diameters estimated by the CDA (D_{TP}). We calculated a percentage difference in the following way:

$$\frac{D_{TP} - D_{GT}}{D_{TP}} \times 100$$

Positive values reflect an overestimation, while negative values are an underestimation of the crater diameter by the algorithm. Any True Positive detections with diameter estimations beyond ±50% were rerecorded as a False Positive. Even though the CDA has identified the correct position of the crater, a diameter estimation beyond ±50% is not accurate nor reliable and therefore not useable for analysis (the True Positive data are available with the supplemental Data Set S2).

Table 3
Recall Values for the "Large Crater Data Set" Degradation Evaluation

| Degradation Class | Diameter (m) | IA-1 | | IA-2 | | IA-3 | | Average |
| | | 66.7° | 66.9° | 51.7° | 56.4° | 41.8° | 48.9° | |
		Near side Mare	Near side Highland	Far side Mare	Far side Highland	Near side Mare	Far side Highland	
A	≥100	0.91	1.00	1.00	1.00	1.00	1.00	0.98
	≥300	–	1.00	–	1.00	–	–	1.00
	≥500	–	1.00	–	–	–	–	1.00
B	≥100	1.00	0.98	1.00	1.00	0.93	0.98	0.98
	≥300	1.00	1.00	1.00	–	1.00	1.00	1.00
	≥500	1.00	1.00	–	–	1.00	–	1.00
C	≥100	0.88	0.80	0.96	0.93	0.89	0.88	0.89
	≥300	0.74	0.74	–	0.70	1.00	0.88	0.81
	≥500	0.75	0.78	–	0.88	–	1.00	0.85

Note. Values are grouped by degradation class (A–C) and lighting condition (IA-1–3; the results for each cell are in Table S2 in Supporting Information S1).

The results show that across both the Small and Large Crater Data Sets, the algorithm overestimates the craters true diameter by ∼15% (Figure 9). Specifically, the overestimation is less across the Large Crater Data Set (100 m to 1 km), with the median sitting at ∼10% (Figure 9b). In comparison, manual mapping of impact craters shows a variability in the crater size measurements of up to ±15% (Robbins et al., 2014). Although the origin of the crater size overestimation by our network is currently unknown, it is unlikely that this comes from the training data set. This will be the subject of future investigation and improvement using other object detection architectures such as YOLOv5.

4. Conclusion

Automated crater mapping is key for quickly and consistently mapping small impact craters across heavily craters surfaces, such as the Moon. Current lunar global impact crater data sets are limited to craters >1–2 km in diameter (e.g., Robbins, 2019), which restricts the ability to gather small-scale information on the physical properties and degradation rates of geological surfaces, as well as the model age of the youngest events occurred on the Moon, such as Copernican impacts. Understanding the spatial densities of small craters can link primary impact events to secondary craters clusters and ejecta (Lagain, Benedix, et al., 2021) and aid in future lunar missions.

By adapting a published image-based CNN algorithm originally trained in the detection of Martian craters (see, Benedix et al., 2020; Lagain, Servis, et al., 2021) for use across the Moon using LRO-NAC images, we demonstrate the versatility of machine learning in planetary mapping. The methodology, technical description, and rationale within this paper showcase the workflow and accuracy of our CDA. We trained a detection model on 43,402 lunar craters identified in NAC images, with an initial transfer learning cycle using a published Martian crater detection model (see, Benedix et al., 2020; Lagain, Benedix, et al., 2021; Lagain, Servis, et al., 2021). To evaluate our model, we ran the CDA across six NAC images at different resolution scales. The evaluation focused on quantifying the effectiveness of the CDA model across crater diameter, lighting conditions, lunar terrains, crater degradation, and diameter estimation.

We show our CDA has good performance metrics across NAC images with lighting conditions/incidence angles between ∼50° and ∼70° regardless of lunar terrain, with an average Recall of 0.93, Precision of 0.7, and a F1 score of 0.79 for impact craters larger than 20 m. From 100 m in diameter, we find that Recall, Precision, and F1 score are 0.86, 0.68, and 0.74. Analyzing images with incidence angles lower than ∼50° leads to a significant decrease in performance. Our network is very effective at detecting fresh and moderately degraded craters but less efficient in detecting very degraded craters, with a loss of 0.1–0.2 for all metrics. Although our analysis shows
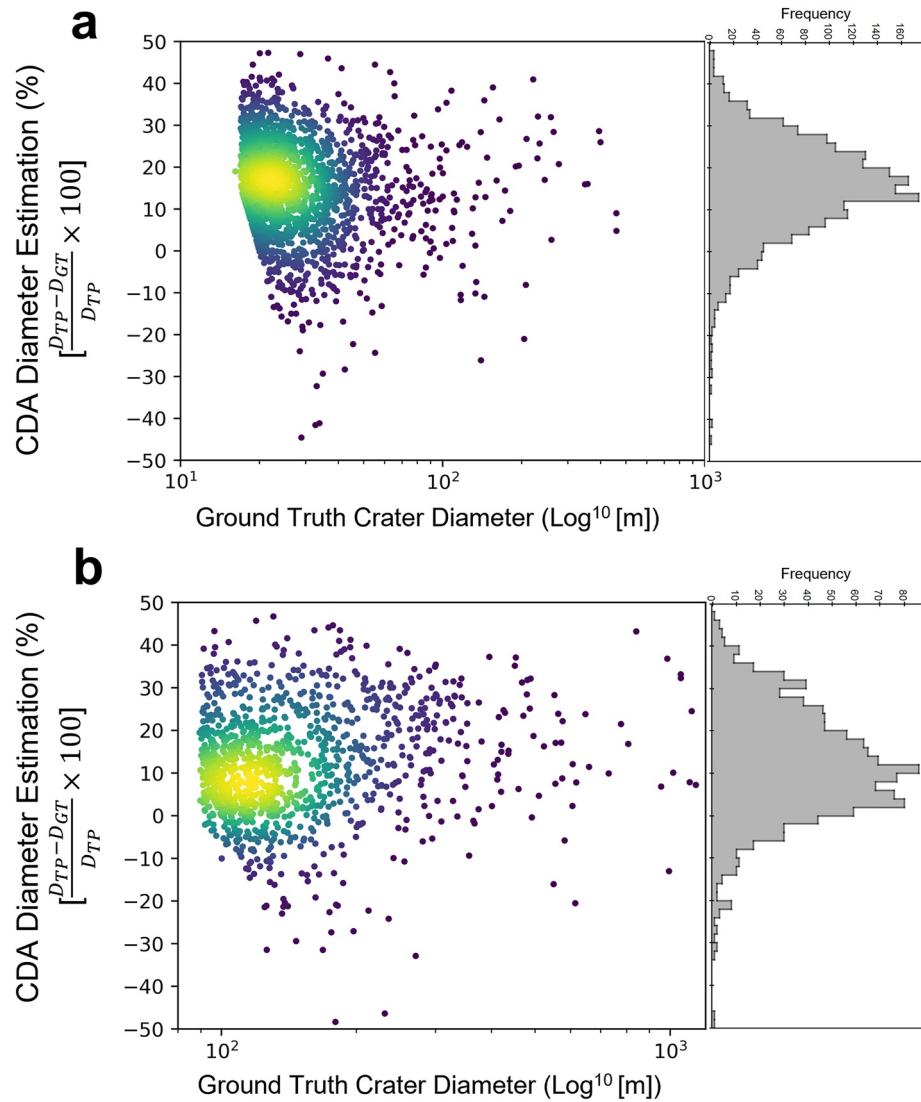
Figure 9. True Positive crater diameter estimation by the Crater Detection Algorithm (CDA) for (a) Small Crater Data Set ($n = 2{,}058$, $D = 20$–500 m), and (b) Large Crater Data Set ($n = 1{,}203$, $D = 100$ m to 1 km). The color coding, via Kernel Density Estimation (KDE), which helps visualize the spatial density of the points relative to each other, yellow indicates more craters and blue indicates fewer craters. The frequency of detection size estimation is computed per increment of 2%. For the vertical axis, D_{TP} refers to the CDA crater diameter of true positive detection and D_{GT} refers to the ground truth crater diameter.

that the CDA displays a 15% overestimation of the crater diameter, we have created the infrastructure needed to process thousands of raw LRO-NAC images, embed the images with geographical information and convert them into a useable file format. Our pipeline allows us to run these images through a trained lunar crater model to generate fast crater detections with acceptable accuracy for detecting fresh and moderately degraded impact craters >20 m on NAC images with incidence angle >50° covering both Mare and Highland terrains.

Data Availability Statement

All the processing steps, publicly available codes, and materials required to reproduce the presented CDA and evaluation are listed here, described throughout the paper, and located within the supplemental materials. The supplementary data are available at Fairweather et al. (2022). The version and implementation of YOLO we used for our Crater Detection Algorithm is available at Ultralytics-yolov3 (2020). The tool we used to label

craters in the ground truth data set (available in the supplementary material, link above) is available at Yolo_Label (2021). The NAC images were downloaded using the LROC website located at The Lunar Reconnaissance Orbiter Narrow-Angle Camera (2022). The United States Geological Survey Integrated Software for Imagers and Spectrometers (ISIS) used to process the downloaded NAC images is available at ISIS3 (2021) or ISIS Documentation (2021).

References

Baldwin, R. B. (1964). Lunar crater counts. *The Astronomical Journal*, *69*, 377. https://doi.org/10.1086/109289

Baldwin, R. B. (1965). Mars: An estimate of the age of its surface. *Science*, *149*(3691), 1498–1499. https://doi.org/10.1126/science.149.3691.1498

Baldwin, R. B. (1987a). On the relative and absolute ages of seven lunar front face basins: I. From viscosity arguments. *Icarus*, *71*(1), 1–18. https://doi.org/10.1016/0019-1035(87)90158-8

Baldwin, R. B. (1987b). On the relative and absolute ages of seven lunar front face basins: II. From the crater counts. *Icarus*, *71*(1), 19–29. https://doi.org/10.1016/0019-1035(87)90159-X

Bart, G. D., Nickerson, R. D., Lawder, M. T., & Melosh, H. J. (2011). Global survey of lunar regolith depths from LROC images. *Icarus*, *215*(2), 485–490. https://doi.org/10.1016/j.icarus.2011.07.017

Benedix, G. K., Lagain, A., Chai, K., Meka, S., Anderson, S., Norman, C., et al. (2020). Deriving surface ages on Mars using automated Crater counting. *Earth and Space Science*, *7*, e2019EA001005. https://doi.org/10.1029/2019EA001005

Bhavsar, H., & Ganatra, A. (2012). A comparative study of training algorithms for supervised machine learning. *International Journal of Soft Computing and Engineering*, *2*(4), 74–81.

Cadogan, P. H. (2020). Automated precision counting of very small craters at lunar landing sites. *Icarus*, *348*, 113822. https://doi.org/10.1016/j.icarus.2020.113822

Crater Analysis Techniques Working Group. (1979). Standard techniques for presentation and analysis of crater size–frequency data. *Icarus*, *37*(2), 467–474. https://doi.org/10.1016/0019-1035(79)90009-5

Daubar, I. J., Atwood-Stone, C., Byrne, S., McEwen, A. S., & Russell, P. S. (2014). The morphology of small fresh craters on Mars and the Moon: Fresh crater d/D on Mars and the Moon. *Journal of Geophysical Research: Planets*, *119*, 2620–2639. https://doi.org/10.1002/2014JE004671

DeLatte, D. M., Crites, S. T., Guttenberg, N., & Yairi, T. (2019). Automated Crater Detection Algorithms from a machine learning perspective in the Convolutional Neural Network era. *Advances in Space Research*, *64*(8), 1615–1628. https://doi.org/10.1016/j.asr.2019.07.017

Di, K., Li, W., Yue, Z., Sun, Y., & Liu, Y. (2014). A machine learning approach to crater detection from topographic data. *Advances in Space Research*, *54*(11), 2419–2429. https://doi.org/10.1016/j.asr.2014.08.018

Di Tommaso, P., Chatzou, M., Floden, E. W., Barja, P. P., Palumbo, E., & Notredame, C. (2017). Nextflow enables reproducible computational workflows. *Nature Biotechnology*, *35*(4), 316–319. https://doi.org/10.1038/nbt.3820

Fairweather, J. H., Lagain, A., Servis, K., Benedix, G. K., Kumar, S. S., & Bland, P. A. (2022). Supporting Material for "Automatic Mapping of Small Lunar Impact Craters Using LROC NAC Images" (Version 1). [dataset]. Zenodo. https://doi.org/10.5281/zenodo.6386231

Fassett, C. I. (2016). Analysis of impact crater populations and the geochronology of planetary surfaces in the inner solar system: Crater populations and surface chronology. *Journal of Geophysical Research: Planets*, *121*, 1900–1926. https://doi.org/10.1002/2016JE005094

Fassett, C. I., & Thomson, B. J. (2014). Crater degradation on the lunar Maria: Topographic diffusion and the rate of erosion on the Moon: Crater degradation on the lunar Maria. *Journal of Geophysical Research: Planets*, *119*, 2255–2271. https://doi.org/10.1002/2014JE004698

Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.

Hartmann, W. K. (1965). Terrestrial and lunar flux of large meteorites in the last two billion years. *Icarus*, *4*(2), 157–165. https://doi.org/10.1016/0019-1035(65)90057-6

Hartmann, W. K. (1966). Martian cratering. *Icarus*, *5*(1–6), 565–576. https://doi.org/10.1016/0019-1035(66)90071-6

Hartmann, W. K. (1971). Martian cratering III: Theory of crater obliteration. *Icarus*, *15*(3), 410–428. https://doi.org/10.1016/0019-1035(71)90119-9

Hartmann, W. K. (1977). Relative crater production rates on planets. *Icarus*, *31*(2), 260–276. https://doi.org/10.1016/0019-1035(77)90037-9

Hashimoto, S., & Mori, H. (2019). *Lunar Crater detection based on grid partition using deep learning*. Paper presented at 2019 IEEE 13th International Symposium on Applied Computational Intelligence and Informatics (SACI) (pp. 75–80). https://doi.org/10.1109/SACI46893.2019.9111474

Head, J. W., Fassett, C. I., Kadish, S. J., Smith, D. E., Zuber, M. T., Neumann, G. A., & Mazarico, E. (2010). Global distribution of large lunar craters: Implications for resurfacing and impactor populations. *Science*, *329*(5998), 1504–1507. https://doi.org/10.1126/science.1195050

Herrick, R. R., Bateman, E. M., Crumpacker, W. G., & Bates, D. (2018). Observations from a global database of impact craters on Mercury with diameters greater than 5 km. *Journal of Geophysical Research: Planets*, *123*, 2089–2109. https://doi.org/10.1029/2017JE005516

Hiesinger, H., Head, J. W., Wolf, U., Jaumann, R., & Neukum, G. (2011). Ages and stratigraphy of lunar mare basalts: A synthesis. In W. A. Ambrose & D. A. Williams (Eds.), *Recent advances and current research issues in lunar stratigraphy*. Geological Society of America. https://doi.org/10.1130/2011.2477(01)

ISIS3. (2021). *[C++]. (Version 3). USGS Astrogeology*. Retrieved from https://github.com/USGS-Astrogeology/ISIS3

ISIS Documentation. (2021). *ISIS: Integrated Software for Imagers and Spectrometers*. Retrieved from https://isis.astrogeology.usgs.gov/User-Docs/index.html

Ivanov, B. A., Neukum, G., Bottke, W. F., & Hartmann, W. K. (2002). The comparison of size–frequency distributions of impact craters and asteroids and the planetary cratering rate. In *Asteroids III* (p. 1). Retrieved from https://www.researchgate.net/publication/253429785

Jia, M., Yue, Z., Di, K., Liu, B., Liu, J., & Michael, G. (2020). A catalogue of impact craters larger than 200 m and surface age analysis in the Chang'e-5 landing area. *Earth and Planetary Science Letters*, *541*, 116272. https://doi.org/10.1016/j.epsl.2020.116272

Jia, Y., Liu, L., Wan, G., & Zhang, C. (2020). *Research progress of lunar impact crater detection*. Paper presented at 2020 International Conference on Computer Science and Management Technology (ICCSMT) (pp. 91–95). https://doi.org/10.1109/ICCSMT51754.2020.00025

Jocher, G. (2021). *Ultralytics YOLOv3 [Python]. Ultralytics*. Retrieved from https://github.com/ultralytics/yolov3

Kneissl, T., van Gasselt, S., & Neukum, G. (2011). Map-projection-independent crater size–frequency determination in GIS environments—New software tool for ArcGIS. *Planetary and Space Science*, *59*(11–12), 1243–1254. https://doi.org/10.1016/j.pss.2010.03.015

Krüger, T., Hergarten, S., & Kenkmann, T. (2018). Deriving morphometric parameters and the simple-to-complex transition diameter from a high-resolution, global database of fresh lunar impact craters (D ≥ ∼3 km). *Journal of Geophysical Research: Planets*, *123*, 2667–2690. https://doi.org/10.1029/2018JE005545

Lagain, A., Benedix, G. K., Servis, K., Baratoux, D., Doucet, L. S., Rajšic, A., et al. (2021). The Tharsis mantle source of depleted shergottites revealed by 90 million impact craters. *Nature Communications*, *12*(1), 6352. https://doi.org/10.1038/s41467-021-26648-3

Lagain, A., Bouley, S., Baratoux, D., Marmo, C., Costard, F., Delaa, O., et al. (2021). Mars crater database: A participative project for the classification of the morphological characteristics of large Martian craters. In W. U. Reimold & C. Koeberl (Eds.), *Large meteorite impacts and planetary evolution VI*. Geological Society of America. https://doi.org/10.1130/2021.2550(29)

Lagain, A., Servis, K., Benedix, G. K., Norman, C., Anderson, S., & Bland, P. A. (2021). Model age derivation of large Martian impact craters, using automatic crater counting methods. *Earth and Space Science*, *8*, e2020EA001598. https://doi.org/10.1029/2020EA001598

Langner, R., & Eickhoff, S. B. (2013). Sustaining attention to simple tasks: A meta-analytic review of the neural mechanisms of vigilant attention. *Psychological Bulletin*, *139*(4), 870–900. https://doi.org/10.1037/a0030694

Lee, C. (2019). Automated crater detection on Mars using deep learning. *Planetary and Space Science*, *170*, 16–28. https://doi.org/10.1016/j.pss.2019.03.008

Lee, C., & Hogan, J. (2021). Automated crater detection with human level performance. *Computers & Geosciences*, *147*, 104645. https://doi.org/10.1016/j.cageo.2020.104645

Liu, Z., Yue, Z., Michael, G. G., Gou, S., Di, K., Sun, S., & Liu, J. (2018). *A global database and analyses of (4) Vesta craters*. Paper presented at 49th Annual Lunar and Planetary Science Conference (Vol. 49, abstract #2083).

Mahanti, P., Robinson, M. S., Thompson, T. J., & Henriksen, M. R. (2018). Small lunar craters at the Apollo 16 and 17 landing sites—Morphology and degradation. *Icarus*, *299*, 475–501. https://doi.org/10.1016/j.icarus.2017.08.018

McEwen, A. S., & Bierhaus, E. B. (2006). The importance of secondary cratering to age constraints on planetary surfaces. *Annual Review of Earth and Planetary Sciences*, *34*(1), 535–567. https://doi.org/10.1146/annurev.earth.34.031405.125018

Melosh, H. J. (1989). *Impact cratering: A geologic process*. Oxford University Press. Retrieved from https://books.google.com.au/books?id=nZwRAQAAIAAJ

Minton, D. A., Fassett, C. I., Hirabayashi, M., Howl, B. A., & Richardson, J. E. (2019). The equilibrium size–frequency distribution of small craters reveals the effects of distal ejecta on lunar landscape morphology. *Icarus*, *326*, 63–87. https://doi.org/10.1016/j.icarus.2019.02.021

Moore, H. J., Boyce, J. M., & Hahn, D. A. (1980). Small impact craters in the lunar regolith—Their morphologies, relative ages, and rates of formation. *The Moon and the Planets*, *23*(2), 231–524.

Nelson, D. M., Koeber, S. D., Daud, K., Robinson, M. S., Watters, T. R., Banks, M. E., & Williams, N. R. (2014). *Mapping lunar Maria extents and lobate scarps using LROC image products*. Paper presented at Lunar and Planetary Science Conference (Vol. 45, abstract #2861).

Neukum, G., Ivanov, B. A., & Hartmann, W. K. (2001). Cratering records in the inner solar system in relation to the lunar reference system. In R. Kallenbach, J. Geiss, & W. K. Hartmann (Eds.), *Chronology and evolution of Mars* (Vol. 12, pp. 55–86). Springer. https://doi.org/10.1007/978-94-017-1035-0_3

Öpik, E. J. (1960). The lunar surface as an impact counter. *Monthly Notices of the Royal Astronomical Society*, *120*(5), 404–411. https://doi.org/10.1093/mnras/120.5.404

Pohn, H. A., & Offield, T. W. (1970). Lunar crater morphology and relative-age determination of lunar geologic units—Part 1. Classification. In *US Geological Survey Prof. Paper* (pp. 153–162).

Povilaitis, R. Z., Robinson, M. S., van der Bogert, C. H., Hiesinger, H., Meyer, H. M., & Ostrach, L. R. (2018). Crater density differences: Exploring regional resurfacing, secondary crater populations, and crater saturation equilibrium on the Moon. *Planetary and Space Science*, *162*, 41–51. https://doi.org/10.1016/j.pss.2017.05.006

Qiao, L., Head, J. W., Xiao, L., Wilson, L., & Dufek, J. D. (2018). The role of substrate characteristics in producing anomalously young crater retention ages in volcanic deposits on the Moon: Morphology, topography, subresolution roughness, and mode of emplacement of the Sosigenes lunar irregular mare patch. *Meteoritics & Planetary Science*, *53*(4), 778–812. https://doi.org/10.1111/maps.13003

Rajšić, A., Miljković, K., Wójcicka, N., Collins, G. S., Onodera, K., Kawamura, T., et al. (2021). Numerical simulations of the Apollo S-IVB artificial impacts on the Moon. *Earth and Space Science*, *8*, e2021EA001887. https://doi.org/10.1029/2021EA001887

Redmon, J., & Farhadi, A. (2018). *YOLOv3: An incremental improvement*. ArXiv Preprint ArXiv Retrieved from https://arxiv.org/abs/1804.02767

Redmon, J., Vivvala, S., Girshick, R., & Farhadi, A. (2016). *You only look once: Unified, real-time object detection*. Paper presented at Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 779–788).

Richardson, M., Malagón, A. A. P., Lebofsky, L. A., Grier, J., Gay, P., Robbins, S. J., & Team, T. C. (2022). The CosmoQuest Moon Mappers Community Science Project: The Effect of Incidence Angle on the Lunar Surface Crater Distribution. https://doi.org/10.48550/ARXIV.2110.13404

Robbins, S. J. (2019). A new global database of lunar impact craters >1–2 km: 1. Crater locations and sizes, comparisons with published databases, and global analysis. *Journal of Geophysical Research: Planets*, *124*, 871–892. https://doi.org/10.1029/2018JE005592

Robbins, S. J., Antonenko, I., Kirchoff, M. R., Chapman, C. R., Fassett, C. I., Herrick, R. R., et al. (2014). The variability of crater identification among expert and community crater analysts. *Icarus*, *234*, 109–131. https://doi.org/10.1016/j.icarus.2014.02.022

Robbins, S. J., & Hynek, B. M. (2012). A new global database of Mars impact craters ≥1 km: 1. Database creation, properties, and parameters. *Journal of Geophysical Research*, *117*, E05004. https://doi.org/10.1029/2011JE003966

Robbins, S. J., Riggs, J. D., Weaver, B. P., Bierhaus, E. B., Chapman, C. R., Kirchoff, M. R., et al. (2018). Revised recommended methods for analyzing crater size–frequency distributions. *Meteoritics & Planetary Science*, *53*(4), 891–931. https://doi.org/10.1111/maps.12990

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

Salamunićcar, G., Lončarić, S., Grumpe, A., & Wöhler, C. (2014). Hybrid method for crater detection based on topography reconstruction from optical images and the new LU78287GT catalogue of Lunar impact craters. *Advances in Space Research*, *53*(12), 1783–1797. https://doi.org/10.1016/j.asr.2013.06.024

Salamunićcar, G., Lončarić, S., & Mazarico, E. (2012). LU60645GT and MA132843GT catalogues of lunar and Martian impact craters developed using a crater shape-based interpolation Crater Detection Algorithm for topography data. *Planetary and Space Science*, *60*(1), 236–247. https://doi.org/10.1016/j.pss.2011.09.003

Salamunićcar, G., Lončarić, S., Pina, P., Bandeira, L., & Saraiva, J. (2011). MA130301GT catalogue of Martian impact craters and advanced evaluation of Crater Detection Algorithms using diverse topography and image datasets. *Planetary and Space Science*, *59*(1), 111–131. https://doi.org/10.1016/j.pss.2010.11.003

Schultz, P., Greeley, R., & Gault, D. (1977). *Interpreting statistics of small lunar craters, Vol. 3, planetary and lunar surfaces*.

See, J. E., Howe, S. R., Warm, J. S., & Dember, W. N. (1995). Meta-analysis of the sensitivity decrement in vigilance. *Psychological Bulletin*, *117*(2), 230–249. https://doi.org/10.1037/0033-2909.117.2.230

Shoemaker, E. M. (1964). The geology of the Moon. *Scientific American*, *211*, 38–47.

Shoemaker, E. M., & Hackman, R. J. (1962). Stratigraphic basis for a lunar time scale. *The Moon*, *14*, 289–300.

Silburt, A., Ali-Dib, M., Zhu, C., Jackson, A., Valencia, D., Kissin, Y., et al. (2019). Lunar crater identification via deep learning. *Icarus*, *317*, 27–38. https://doi.org/10.1016/j.icarus.2018.06.022

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): Lola initial observations. *Geophysical Research Letters*, *37*, L18204. https://doi.org/10.1029/2010GL043751

Soderblom, L. A. (1970). A model for small-impact erosion applied to the lunar surface. *Journal of Geophysical Research*, *75*(14), 2655–2661. https://doi.org/10.1029/JB075i014p02655

Speyerer, E. J., Robinson, M. S., & Denevi, B. W., & LROC Team. (2011). *Lunar Reconnaissance Orbiter Camera morphological map of the Moon*. Paper presented at Lunar and Planetary Science Conference (Vol. 42, abstract #2387). Retrieved from https://www.lpi.usra.edu/meetings/lpsc2011/pdf/2387.pdf

Stadermann, A. C., Zanetti, M. R., Jolliff, B. L., Hiesinger, H., van der Bogert, C. H., & Hamilton, C. W. (2018). The age of lunar mare basalts south of the Aristarchus Plateau and effects of secondary craters formed by the Aristarchus event. *Icarus*, *309*, 45–60. https://doi.org/10.1016/j.icarus.2018.02.030

Stepinski, T. F., Ding, W., & Vilalta, R. (2012). Detecting impact craters in planetary images using machine learning. In *Intelligent data analysis for real-life applications: Theory and practice* (pp. 146–159). IGI Global.

Still, M. (2006). *The definitive guide to ImageMagick*. Apress. Retrieved from https://link.springer.com/content/pdf/10.1007/978-1-4302-0112-0.pdf

Stopar, J. D., Robinson, M. S., Barnouin, O. S., McEwen, A. S., Speyerer, E. J., Henriksen, M. R., & Sutton, S. S. (2017). Relative depths of simple craters and the nature of the lunar regolith. *Icarus*, *298*, 34–48. https://doi.org/10.1016/j.icarus.2017.05.022

Stopar, J. D., Robinson, M. S., Speyerer, E. J., Burns, K., Gengl, H., & Team, L. (2012). *Regolith characterization using LROC NAC Digital Elevation Models of small lunar craters*. Paper presented at Lunar and Planetary Science Conference (Vol. 43, abstract #2387).

The Lunar Reconnaissance Orbiter Narrow-Angle Camera. (2022). *LROC Image Search*. Retrieved from http://wms.lroc.asu.edu/lroc/search

Trask, N. J. (1971). Geologic comparison of mare materials in the lunar equatorial belt, including Apollo 11 and Apollo 12 landing sites. *U. S. Geological Survey Professional Paper*, *750D*, 138–148.

Ultralytics-yolov3. (2020). *(Python). [Code] Ultralytics*. (Original work published 2018). Retrieved from https://github.com/ultralytics/yolov3

Urbach, E. R., & Stepinski, T. F. (2009). Automatic detection of sub-km craters in high resolution planetary images. *Planetary and Space Science*, *57*(7), 880–887. https://doi.org/10.1016/j.pss.2009.03.009

van der Bogert, C. H., Hiesinger, H., Dundas, C. M., Krüger, T., McEwen, A. S., Zanetti, M., & Robinson, M. S. (2017). Origin of discrepancies between crater size–frequency distributions of coeval lunar geologic units via target property contrasts. *Icarus*, *298*, 49–63. https://doi.org/10.1016/j.icarus.2016.11.040

Wang, Y., Wu, B., Xue, H., Li, X., & Ma, J. (2021). An improved global catalog of lunar impact craters (≥1 km) with 3D morphometric information and updates on global crater analysis. *Journal of Geophysical Research: Planets*, *126*, e2020JE006728. https://doi.org/10.1029/2020JE006728

Wang, Y., Xie, M., Xiao, Z., & Cui, J. (2020). The minimum confidence limit for diameters in crater counts. *Icarus*, *341*, 113645. https://doi.org/10.1016/j.icarus.2020.113645

Wilcox, B. B., Robinson, M. S., Thomas, P. C., & Hawke, B. R. (2005). Constraints on the depth and variability of the lunar regolith. *Meteoritics & Planetary Science*, *40*(5), 695–710. https://doi.org/10.1111/j.1945-5100.2005.tb00974.x

Wilhelms, D. E. (1987). *The geologic history of the Moon* (USGS Numbered Series No. 1348; Professional Paper, 337). United States Geological Survey. https://doi.org/10.3133/pp1348

Williams, J. P., van der Bogert, C. H., Pathare, A. V., Michael, G. G., Kirchoff, M. R., & Hiesinger, H. (2018). Dating very young planetary surfaces from crater statistics: A review of issues and challenges. *Meteoritics & Planetary Science*, *53*(4), 554–582. https://doi.org/10.1111/maps.12924

Xiao, Z., & Strom, R. G. (2012). Problems determining relative and absolute ages using the small crater population. *Icarus*, *220*(1), 254–267. https://doi.org/10.1016/j.icarus.2012.05.012

Xiao, Z., & Werner, S. C. (2015). Size–frequency distribution of crater populations in equilibrium on the Moon: Crater equilibrium on the Moon. *Journal of Geophysical Research: Planets*, *120*, 2277–2292. https://doi.org/10.1002/2015JE004860

Yang, C., Zhao, H., Bruzzone, L., Benediktsson, J. A., Liang, Y., Liu, B., et al. (2020). Lunar impact crater identification and age estimation with Chang'E data by deep and transfer learning. *Nature Communications*, *11*(1), 6358. https://doi.org/10.1038/s41467-020-20215-y

Yolo_Label. (2021). *(Version 1.1.1) [Code]*. Retrieved from https://github.com/developer0hye/Yolo_Label

Young, J. (1940). A statistical investigation of diameter and distribution of lunar craters. *Journal of the British Astronomical Association*, *50*(9), 309–326.

Zeilnhofer, M. (2020). *A global analysis of impact craters on Ceres* (ProQuest Dissertations & Theses Global, pp. 1–251, proquest: 27963349). Retrieved from https://search.proquest.com/openview/4e9a383f80b524bf68f755ff3649d02d

Zhang, Y. (2010). *New advances in machine learning*. Intech. https://doi.org/10.5772/225

References From the Supporting Information

Acton, C. H. (1996). Ancillary data services of NASA's navigation and ancillary information facility. *Planetary and Space Science*, *44*(1), 65–70. https://doi.org/10.1016/0032-0633(95)00107-7

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

GDAL Documentation. (2021). *GDAL*. Retrieved from https://gdal.org/

McMahon, S. K. (1996). Overview of the planetary data system. *Planetary and Space Science*, *44*(1), 3–12. https://doi.org/10.1016/0032-0633(95)00101-8

Scholten, F., Oberst, J., Matz, K.-D., Roatsch, T., Wählisch, M., Speyerer, E. J., & Robinson, M. S. (2012). GLD100: The near-global lunar 100 m raster DTM from LROC WAC stereo image data: The 100 M RASTER DTM GLD100. *Journal of Geophysical Research*, *117*, E00H17. https://doi.org/10.1029/2011JE003926