EXPLORING THE MOON WITH AUTOMATED FEATURE DETECTION. R. V. Wagner, E. J. Speyerer, and M. S. Robinson. School of Earth and Space Exploration, Arizona State University, Tempe, AZ 85287-3603.

Introduction: The Lunar Reconnaissance Orbiter Camera (LROC) consists of a single Wide Angle Camera (WAC) and twin Narrow Angle Cameras (NACs) that provide multispectral and high-resolution imaging, respectively [1]. The NACs are capable of acquiring panchromatic images at 0.5 m/pixel from an altitude of 50 km. A typical NAC image consists of 5064 samples and 52224 lines, resulting in over 500 megapixels of image data in each observation pair. As of 15 March 2017, the NACs have collected over 650,000 image pairs of illuminated terrain (>1.3 million individual images) covering most of the Moon.

Here we present a feature detection tool and an automated change detection tool developed by the LROC team. These tools (*PitScan* and *CRISP*) enable the discovery of lunar pits [2] as well as a measurement of the contemporary impact flux and the rate of regolith overturn [3].

Finding Lunar Pits with *PitScan*: Lunar pits are deep, vertical-walled collapse features, generally <100 m in diameter, and usually have an inward-sloping rim. So far, we have located over 300 pits across the lunar surface. Their small size and rarity make pits a prime candidate for automated searching, and in fact most of the known pits were found using *PitScan* (previously reported in [2]).

Theory. Since the majority of slopes on the Moon are below the angle of repose (~36°) [4], very few features cast shadows when the Sun is within ~54° of the zenith. Pits, with their vertical or near-vertical walls, do cast shadows (as do large boulders, and small outcrops on steep crater slopes). Thus, a catalog of shadows in a high-Sun image should contain any pits in the area. All that is needed is to filter out shadows in pits from shadows of non-pit features.

Implementation. PitScan was developed to locate all shadows larger than 15 pixels across while excluding those features that are most likely to be boulders. All remaining features are saved as small image clippings for a human analyst to check manually. *PitScan* can complete a search of a single 250 megapixel NAC image in thirty seconds.

To find shadows, *PitScan* uses an empirically derived equation to calculate a cutoff value for "shadowed" pixels. The formula for this cutoff value [2] was determined by manual inspection of pits in several dozen calibrated images with various Sun elevations.

Once all the continuous blobs of pixels with I/F values below the cutoff that are at least 15 pixels across have been located, *PitScan* extracts a profile across each blob parallel to the solar azimuth, extending approximately 30 pixels beyond the bounds of the shadow (Fig. 1). If the average I/F value on the up-Sun side of the shadow is greater than 0.9× the average value on the down-Sun side, the feature is assumed to be a rock, and discarded. For the remaining cases, *PitScan* saves a 300 × 300 pixel clipping for human review.

In cases where more than 50 potential pits are found in an image, *PitScan* instead saves a preview of the entire image with potential pit locations marked, so a human can check for known patterns of false positives (such as outcrops on crater walls).

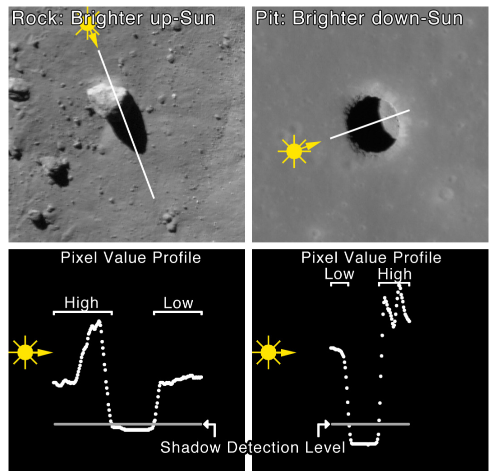

Fig. 1: *Top*: Profile lines across detected shadows for rock (*left*) and pit (*right*). *Bottom*: Pixel values along the profile.

Results. Excluding the feature-rich images, the algorithm currently generates ~150 false positives for each successful pit identification. We consider this an acceptable level of false positives, as an experienced analyst can evaluate most image clippings in less than a second, and it only takes a few hours to check the results for six months of NAC images. The false negative rate can be greatly reduced by adjusting the ratio used to detect boulders: If the down-Sun side is at least 1.1x the up-Sun side, 40% of false negatives are excluded, and 85% of true positives are retained.

In a sample of all images of known pits with pixel scales such that the pit is at least 30 pixels across, *PitScan* only detected 45% of the expected pits. Detection was better for non-impact-melt pits, with 86% of expected detections made, although there may be a sampling bias, as 13 of the 16 known non-impact-melt pits were originally found using *PitScan*, while many impact melt pits were found by manual search near impact melt pits identified by *PitScan*.

Due to the limit on valid Sun elevations, *PitScan* can only search the region within ~50° latitude of the equator (77% of the Moon). To date, the NAC has acquired 307,824 images within the described incidence angle constraint covering 76% of the searchable area.

Finding New Impacts with *CRISP*: With the aid of the Lunar Reconnaissance Orbiter (LRO) extended mission, LROC is collecting repeat images of the same surface area under the same lighting conditions. These temporal image pairs ("before" and "after" images) enable the identification of surface changes over time. To date, we have identified over 200 newly formed impact craters, over 75,000 smaller primary and secondary surface changes, and tens of recent mass wasting events.

Theory. When an impact or mass-wasting event occurs, the local surface reflectance is modified due to changes in roughness, exposed regolith composition, and/or optical maturity. Using temporal image pairs with nearly identical lighting and limited shadows (i.e. incidence angle < 50° and phase difference < 3°), registered images are compared for reflectance and textural changes using an in-house developed tool, Change Recognition using Images with Similar Phase (*CRISP*).

Implementation. CRISP is split into three functions: image registration, pixel- and area-based change-detection filtering, and image segmentation. *CRISP* starts with the PDS formatted and archived Experiment Data Records (EDRs) for each NAC observation that constitutes the temporal pair. Before inspecting the images for changes, the images are calibrated and rectified into a common reference frame using the spacecraft ephemeris and orientation information. To further improve the image registration, such that each surface feature in the before image is matched to the same feature in the after image, the image pairs are registered using an automatic co-registration tool called *coreg* in Integrated Software for Imagers and Spectrometers (ISIS). This approach aligns the images with sub-pixel accuracy enabling pixel-to-pixel comparisons in the temporal pair.

Once registered, an image ratio is created by dividing the after image by the before image (Fig. 2). Since the lighting and viewing geometries are similar and the images are registered within a pixel, the values in the ratio image are near one except in cases where a reflectance change has occurred.

To limit false positives due to slight photometric variations and low signal to noise ratio in shadowed areas, a second change-detection filter is applied. This filter identifies changes in local reflectance patterns in an n × n pixel area (n=11 for current searches of LROC images) using a normalized 2D cross-correlation (NCC) factor. For an area with no detectable surface changes, the corresponding area in the NCC image contains values near one. As the n × n pixel NCC filter encounters a significant texture difference between the "before" and "after" image, the recorded values in the NCC deviate from unity.

Together, the ratio image and NCC image are used to identify surface changes in the temporal image pairs.

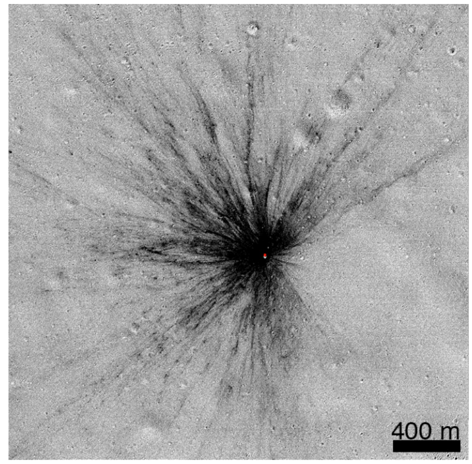

Fig. 2: Example ratio (after/before) image showing the result of a new impact that created a 12 m crater.

The cutoff threshold values were empirically derived using the surface changes identified through manual scanning of 138 NAC temporal pairs [5]. For each detected change identified by *CRISP*, a 200 × 200 pixel cutout is generated for the "before" image, "after" image, ratio image, and change mask. These full-resolution sub-images are combined into animated gif image that is manually inspected by a human analyst and classified.

Results: As of 1 April 2017, the LROC team has systemically scanned and classified 18,516 NAC temporal pairs. The *CRISP* algorithm identified 138,264 potential surface changes (anomalies) in the temporal pairs. Manual classification of these anomalies confirmed 57% were true, distinguishable surface changes. In addition, 19% of the anomalies contained a possible surface change. These possible surface changes were either too small to be definitively confirmed or the images lacked the required signal to adequately confirm the surface change. The remaining 24% of the anomalies were falsely identified by *CRISP* due to resolution difference between the image pair, small lighting differences, and small co-registration errors.

Conclusion: Additional observations acquired during the LRO Cornerstone Mission and future extended missions will enable the LROC team to increase the NAC's spatial and temporal coverage. *PitScan* and *CRISP* will continued to be used to locate more pits as well as refine the contemporary cratering and gardening rate for the Moon as more surface changes are identified. Future efforts will also encompass testing these algorithms on other lunar and planetary datasets to enable further mapping of these features.

References: [1] Robinson, M.S. (2010), *Space Sci. Rev.* 129, 391-419. [2] Wagner, R. V. and Robinson, M. S. (2014), *Icarus*, 237C, 52–60. [3] Speyerer et al. (2016), *Nature*, 538, 215-218. [4] Wagner, R.V. et al. (2013) 44th LPSC, #2924. [5] Thompson et al. (2014) 45th LPSC, #2769.