AUTOMATED BOULDER COUNTING: DEEP LEARNING FOR BOULDER DETECTION AND HEIGHT ESTIMATION. A. H. Stoken[1,2], A. W. Britton[1,2], A. A. Turner[1,2], M. A. Rubio[1,3], M. D. Lambert[1,2], [1]NASA-JSC Houston, TX; [2]Jacobs Inc., Houston, TX; [3]GeoControl Systems Inc., Houston, TX (alex.h.stoken@nasa.gov, andrew.w.britton@nasa.gov, amber.a.turner@nasa.gov, minna.a.rubio@nasa.gov, mark.d.lambert@nasa.gov)

Introduction: For upcoming Artemis missions, conducting hazard assessments across potential landing sites is instrumental in guiding site selection and extra-vehicular activity (EVA) traverse planning. Boulders are an important part of this hazard assessment. Characterizing boulder size, frequency, and spatial distribution across large regions of interest (ROIs) necessitates the development of automated tools that can increase the accuracy and speed of surface hazard detection. Many lunar boulders are difficult to identify and measure in satellite imagery as their extents are unresolvable at the available image resolution. In many cases, boulders are subpixel in diameter, making them nearly undetectable.

In the South Polar region of the Moon, a low Sun elevation angle causes even small boulders to cast long, discernable shadows. In this region, instead of detecting boulders themselves, we can search for boulder shadows. We train a convolutional neural network (CNN) to find boulder shadows in imagery and use the resulting detections to estimate shadow length and compute boulder height.

Data: High resolution data from the Narrow Angle Camera (NAC) instrument aboard NASA's Lunar Reconnaissance Orbiter (LRO) was sourced from the LROC data node at Arizona State University [1] to conduct boulder shadow detection.

Data Specification. Twenty-three NAC frames near the lunar south pole were chosen based on their resolution, incidence angle, and proximity near Connecting Ridge [2]. The finalized NAC frames have a resolution range of 0.7 m to ~ 1.2 m and incidence angle range of ~87° to ~90°. The NAC frames were sliced into 512x512-pixel grayscale images ("tiles") before analysis.

Labeling Training Data. Lunar geospatial scientists analyzed 11,008 NAC frame tiles, identifying individual boulders, boulder clusters, and boulder fields. Boulders were manually identified, interpreted, and labeled [3] with respect to visual and topographical characteristics in-line with the approach of [4]. The geologists used bright sunward facing pixels and other geologic contextual clues (distance to crater rims, ejecta deposits) to identify boulders in the NAC frame tiles.

Methods: We developed a multi-stage processing pipeline for boulder detection and height estimation.

Boulder Shadow Detection. We trained a CNN object detection model on the expert labeled dataset to detect bounding boxes around boulders and the full extent of their associated shadows. We fine-tuned an ImageNet-pretrained YOLOv5s6 model [5]. A round of model-assisted labeling was used to augment the original training set with additional boulder shadows that were originally undetected by human labelers. At inference time, we use SAHI [6] for tiled inference — we predict on tiles extracted from the input image as well as the full image and combine these predictions — giving the model the opportunity to find smaller boulders and shadows.

Azimuth Filtering and Height Estimation. For each shadow bounding box detection, we binarize the area with a threshold value of half the mean pixel brightness in the detection and extract the dark regions. We estimate where a shadow could be by drawing a line from the brightest point of the detection — in most cases the boulder — away from the Sun along the solar azimuth angle. A shadow is considered real where this shadow hypothesis and binarized dark regions overlap. If there is no overlap, we remove the detection from the set of identified boulders, thus filtering out some false positives from the initial detection model.

From the length of the boulder shadow and the solar elevation angle, we use the formula from [4] to calculate boulder height:

$$h = \frac{l}{\tan i} - h'$$

where l is the measured length of the boulder's shadow, i is the solar incidence angle, and h' is the elevation difference between the base of the boulder and the tip of the shadow.

Results: We measured our boulder detection model performance on a held-out validation set of 138 tiles containing 1,261 boulders. Table 1 lists metrics from the final model as well as various ablations. Model training time on an NVIDIA RTX A4000 laptop GPU is approximately 4 hours. Inference speed ranges from 330-593 ms per tile.

Shadow Detection Performance. All results are listed for Intersection over Union threshold of 0.6. Precision and Recall values are averages computed over confidence values from 0.5 to 0.95 in 0.05-point intervals.

Method	Precision	Recall
Full Image Inference	0.306	0.427
Tile-only Inference	0.467	0.543
SAHI (Tile+ Full Inference)	0.497	0.604
w/o Model-Assisted Labels	**0.713**	**0.466**
with Model-Assisted Labels	**0.763**	**0.534**

*Table 1: Average Precision and Average Recall. Values computed from COCO and **YOLO** protocols differ, thus are presented separately. Higher values are better.*

Boulder Spatial and Height Distributions. Figures 1 and 2 show examples of boulder/shadow extraction from detections as well as the spatial distribution of detected boulders in a sample NAC tile.

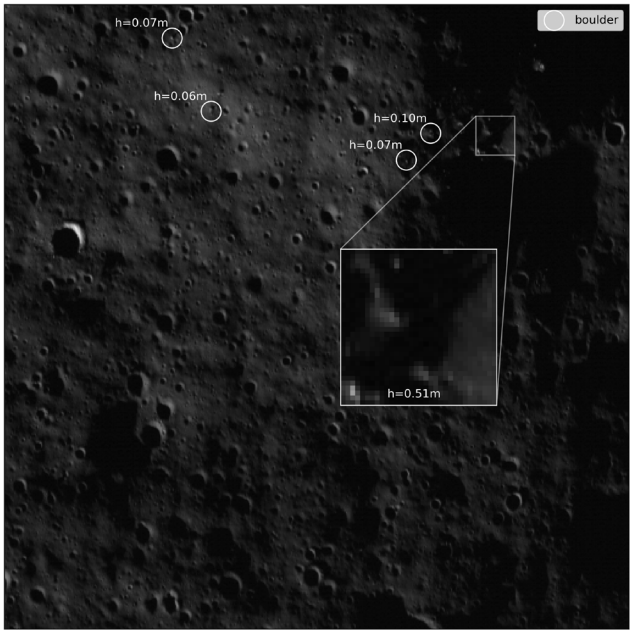

Figure 1: Spatial distribution and heights of boulders in a tile from NAC frame M112504292L. The numbers represent the height in meters of a boulder detected in that position.

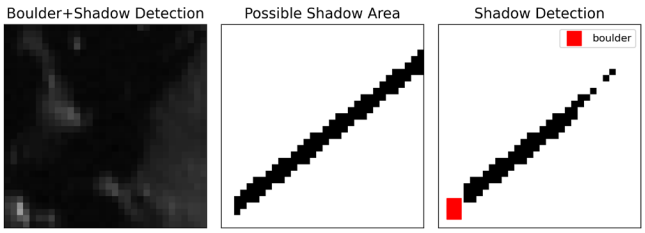

Figure 2: (left) The bounding box containing the boulder and its shadow as detected by the CNN. (middle) The shadow hypothesis — where the shadow could be in the image given the predicted boulder location and the sun azimuth angle. (right) The overlapping region, the length of which is measured to calculate shadow length. Shadow length (l) = 27 m, boulder height (h) = 0.8 m.

Discussion: Generating ground truth training data posed difficulties for lunar geospatial scientists. There are three primary sources of error that arose during the manual NAC frame analysis. The highest contributing factor is the limited spatial resolution, since the image's quality directly affects the ability to confidently assess surface hazards. Secondly, viewing geometry obstructs accurate discernment of boulder frequency and dimensions in highly clustered areas. Thirdly, low Sun elevation angle causes long shadows to be cast by other morphological features such as mounds and crater rims, obscuring smaller features and inhibiting the ability to precisely delineate a boulder and its shadow.

In early iterations of our model, one of the most common misidentifications was labeling small craters as boulders. Craters share some similar visual signatures to boulders but have shadows in the opposite direction, so they can be removed from our boulders lists in the azimuth filtering step. Future work should focus on disambiguating the features at the model stage.

Our method of shadow length estimation has two main sources of error. The first is from the CNN model detection. If the detected region does not include the full extent of the shadow, then our length estimation step does not have enough context to find the shadow tip. When measuring the shadow, we then choose a low threshold for determining the shadow hypothesis to ensure we have as wide a swath of shadow region as possible to intersect with the solar azimuth line, and to ensure that shadows with less contrast are detectable. This choice may lead to overly elongated shadows that then overestimate the height of boulders.

Conclusion: We present a machine learning based method for fast initial hazard assessment. Using a convolutional neural network trained on expertly labeled lunar boulder and shadow data, along with spatially-motivated postprocessing, we show the ability to detect boulders in NAC imagery up to the subpixel level. This method can be a value add to lunar site and EVA planning during Artemis and other future missions.

Acknowledgments: This work was funded as part of the 2022 Jacobs SABRE Grant program on the JETS Contract. We wish to thank those who made the funding possible and saw the potential in our project.

References: [1] Robinson, M. S. (2010). Lunar Reconnaissance Orbiter Camera Experimental Data Record. *LRO-L-LROC-2-EDR-V1. 0*. [2] Gläser, P., et al. (2018) *PSS* 162: 170-178 [3] Lin, T. (2015) *LabelImg. GitCode*, github.com/tzutalin/labelImg. [4] Watkins, R. N., et al. (2019) *JGR: Planets 124.11: 2754-2771*. [5] Jocher, G., et al. (2022). ultralytics/yolov5:YOLOv5(v7.0) doi: 10.5281/zenodo.7347926. [6] Akyon, F. C., et al. (2021). SAHI doi: 10.5281/zenodo.5718950.