Automatic Lunar Crater Detection from Optical Images and Elevation Maps

Susanna Maria Baby
susmaria@stanford.edu

Mara Finkelstein
mfinkels@stanford.edu

Hugo Kitano
hkitano@stanford.edu

Data

We used a variety of data to both train our models and obtain predictions from them.

- **Optical images** of the moon's surface from Lunar Reconnaissance Orbiter Narrow Angle Cameras of size 600×400 pixels with a resolution of 1 meter per pixel, along with an accompanying crater catalog from the Emami lab at the University of Nevada-Reno. Optical images have high resolution, but also have shadows that can interfere with detection.
- **Digital Elevation Models** (DEM's) from the LRO-Kaguya merged DEM, downsampled to 118 meters per pixel. These were accompanied by two crater catalogs, the LROC crater dataset for craters < 20 km in diameter, and the Head et al catalogue for craters > 20 km in diameter. DEM's are not affected by lighting changes, but are expensive to compute at high resolution, and are thus often downsampled.
- To analyze how our models transfer to other planetary bodies, we used DEM's of **Mercury** (from the Mercury MESSENGER Global DEM) and **Mars** (from the Mars MGS MOLA DEM) to obtain crater predictions from models trained on lunar craters. For prediction of craters on Mars, we also used the Salamuniccar MA132843GT crater catalogue as a ground-truth to evaluate our accuracy.
- The most critical issue with this project is that the **ground-truth catalogs are incomplete**. They are all manually recorded by experts, but they do miss craters, and there are sometimes disagreements between them.

Models

We built a variety of tools to analyze our lunar data.

- A **template-matching pipeline**, built from scratch using OpenCV, that adds missing labels to our images before training our models.
- **Faster R-CNN**, a common object-detection model that inputs an image into a convolutional neural network, uses a region proposal network (RPN) that scans the resulting feature map for regions that may have objects in them, then outputs bounding boxes around its predictions. We adapted Faster-RCNN for our project so it detects craters from optical images.
  - ➢ We ran Faster R-CNN as a baseline, and also ran it using template-generated labels.

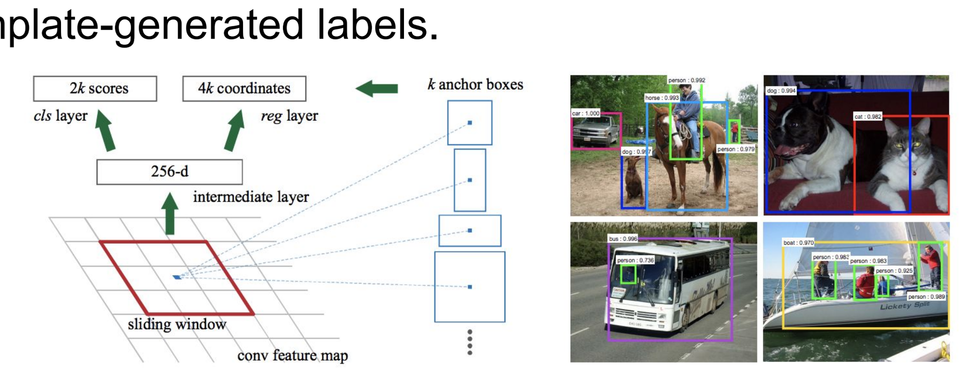

From [3]

- **DeepMoon** is a lunar crater detection model consisting of a convolutional neural network adapted from the UNET architecture. For the first half of its layers, it downsamples, and then upsamples for the second half, all while maintaining short-cut connections between the layers. Finally, it uses template matching to extract the craters from the target masks.
  - ➢ We altered DeepMoon to use dilated convolutions to better identify large craters, which often escaped DeepMoon detection due to being larger than the receptive field size.

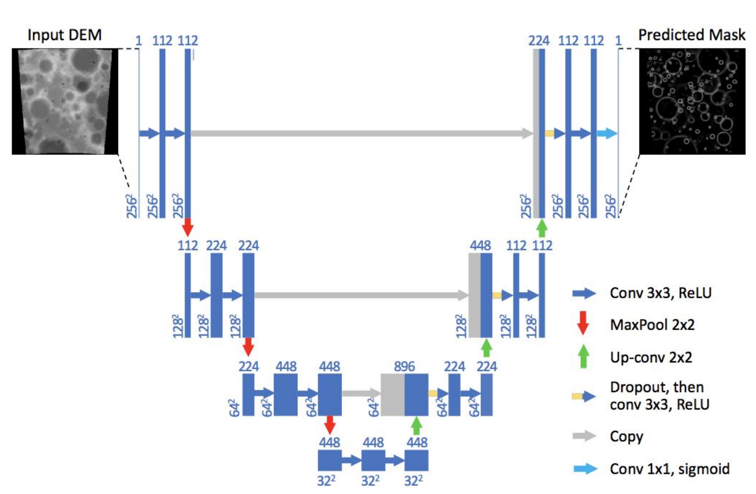

From [1]

Project Summary

Our final project approaches the problem of automatic crater detection using a variety of models and data sources. We trained two types of models, an extension of Faster R-CNN, a common object detection method, and DeepMoon, a state-of-the-art lunar crater detection model, to identify craters and compare them to ground-truth crater catalogs. We trained our models on both optical images and digital elevation models (DEM's) of the Moon's surface. Lastly, we saw that these models were able to detect craters on other planetary bodies (Mercury and Mars) as well.

References

- [1] Silburt, Ari, et al. "DeepMoon: Convolutional neural network trainer to identify moon craters." *Astrophysics Source Code Library* (2018).
- [2] Emami, Ebrahim, et al. "Automatic crater detection using convex grouping and convolutional neural networks." *International Symposium on Visual Computing*. Springer, Cham, 2015.
- [3] Ren, Shaoqing & He, Kaiming & Girshick, Ross & Sun, Jian. (2015). Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. IEEE Transactions on Pattern Analysis and Machine Intelligence. 39.
- [4] Salih, Atheer L., et al. "Automatic Crater Detection and Implications for Surface Age Estimation." *European Planetary Science Congress*. Vol. 9. 2014.

Experiments and Results

- **Faster-RCNN**: Crater discovery using convolutional neural network for object detection

Test set results

Original train and test set labels

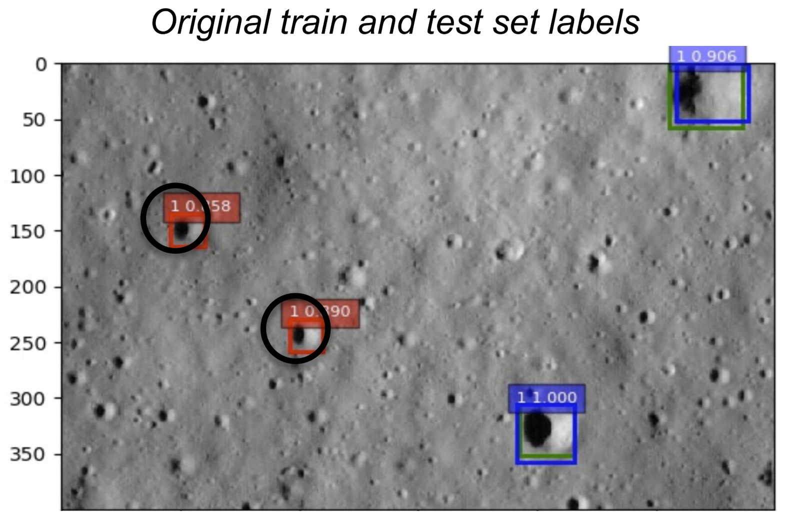

Faster-RCNN test set predictions on model trained without template-generated labels. Both "false positive" predictions (circled in black) are actually true craters.

Labels augmented with template-generated matches

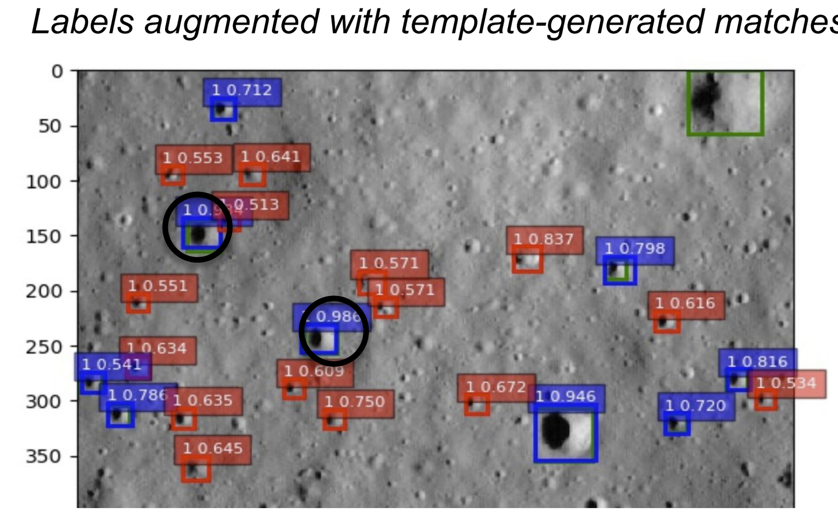

Both of the original false positives are labeled during template matching, and are successfully predicted by the model, but this time as true positives. Note that this model also discovers several new "false positives" now, which are all true craters as well.

- Faster-RCNN Model hyperparameters:
  - Learning rate: 0.001
  - RPN_ANCHOR_SCALES = [8, 16, 24]*ratios (determines effective "receptive field size")
  - IOU thresholds for RPN:
    - RPN_POSITIVE_OVERLAP = 0.5
    - RPN_NEGATIVE_OVERLAP = 0.3
    - Example with IOU between 0.3 and 0.5 not considered during training
  - NMS threshold of 0.7 for filtering regions of interest produced by region proposal network
  - Batch size: 1 image, 128 regions of interest

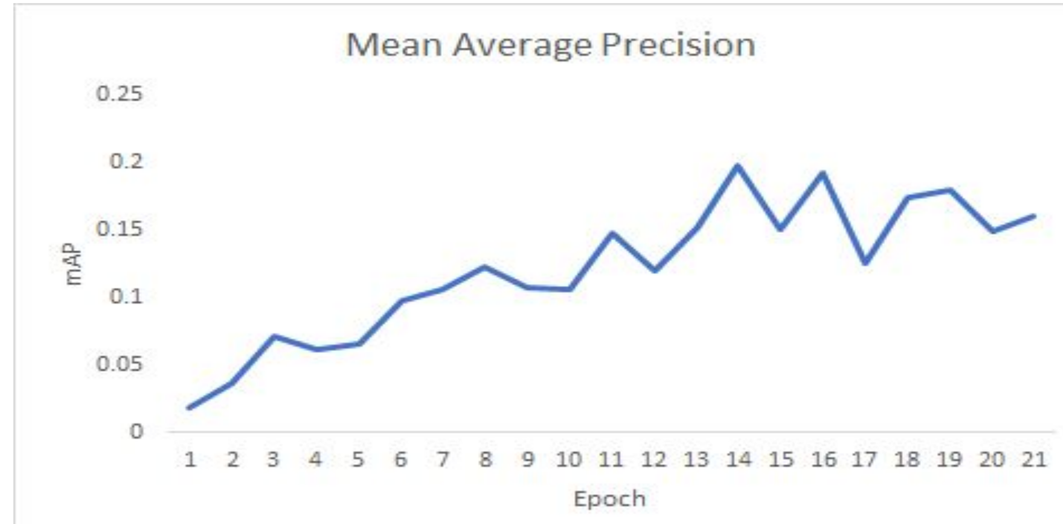

- **DeepMoon:** UNet architecture with crater-specific post-processing

batch loss on validation data for training with and without dilated convolutions

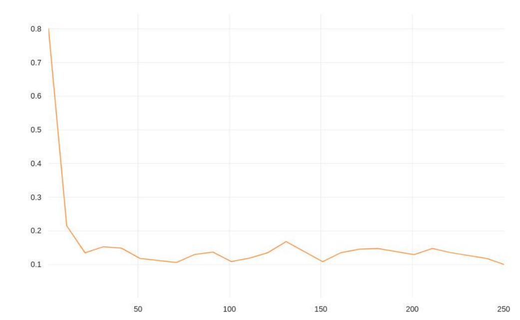

with dilation

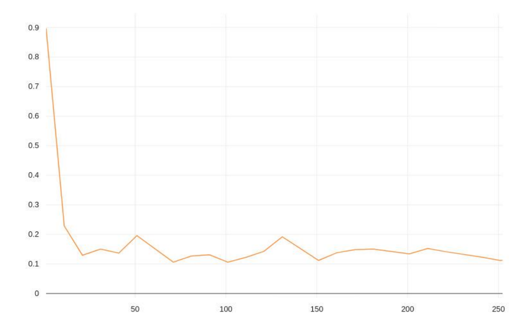

without dilation

The model with dilated convolution layers produces a smoother curve with lower loss in the initial training phase.

- DeepMoon Model hyperparameters:
  - Learning rate: 0.0001
  - Number of filters : 112
  - Cropped image size : 256 x 256
  - Weight initialization : 'he_normal'
  - Weight regularization 'lambda' : 1e-6
  - Dropout fraction : 0.15
  - Batch size: 8

- **Transfer learning** to other astronomical bodies: Mercury and Mars

Failure

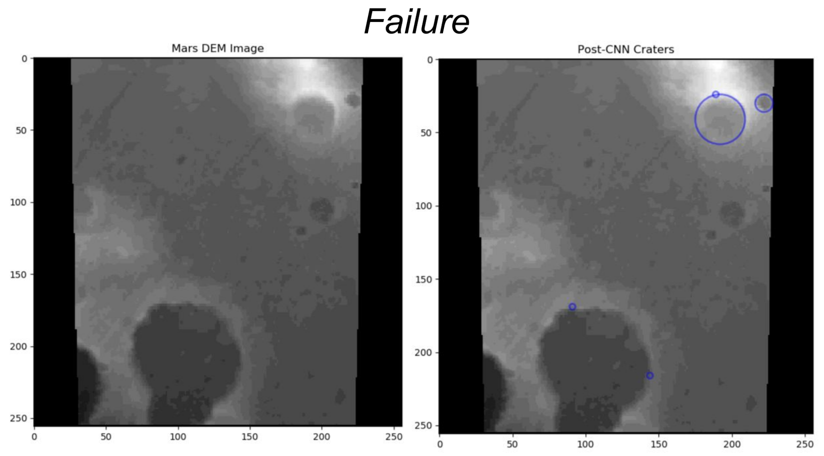

Partial success

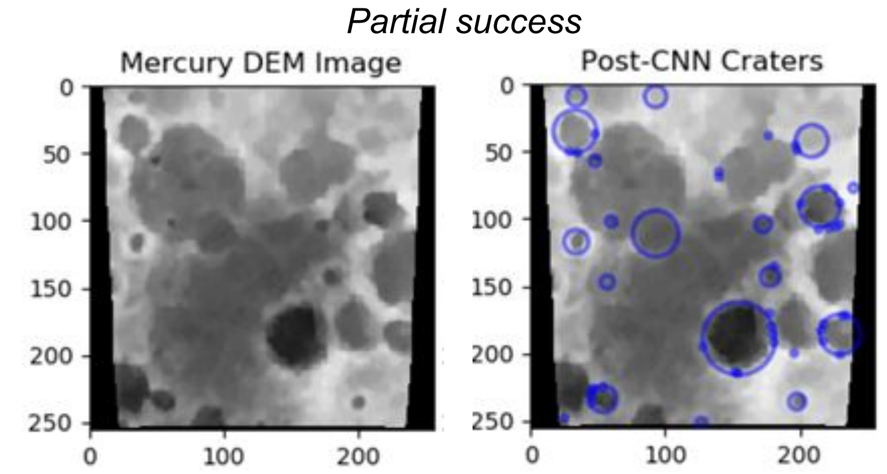

Excellent performance

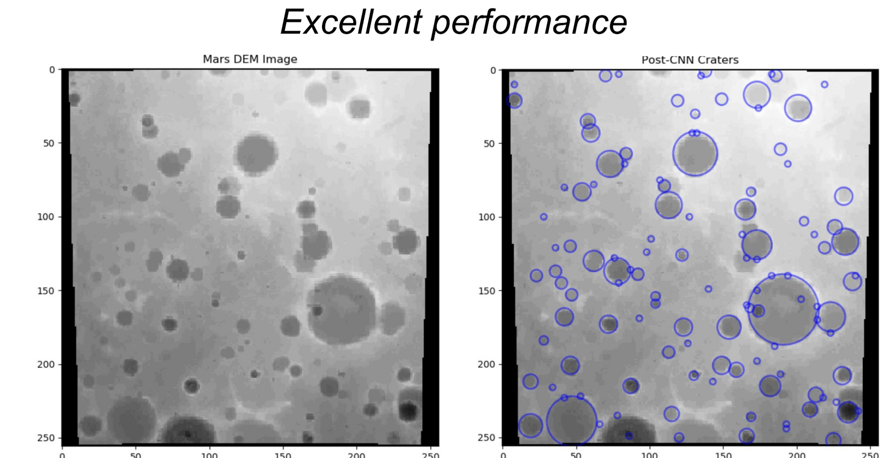

Evaluation

Summary of test set results	No filtering		Filtering for craters with diam. > 20 pixels	
	Precision	Recall	Precision	Recall
RCNN baseline	0.36	0.19	0.41	0.48
RCNN with template-generated labels	0.58	0.22	0.61	0.62
DeepMoon baseline	0.84	0.51	0.33	0.66
DeepMoon with dilated conv	0.85	0.50	0.34	0.64
DeepMoon on optical images	0.64	0.36	0.46	0.59
DeepMoon with dilated conv on optical images	0.67	0.32	0.50	0.55
DeepMoon baseline on Mars DEMs	0.93	0.36	0.34	0.62

Conclusion and Next Steps

- Faster-RCNN and DeepMoon perform well on mid-sized craters. Both need higher-resolution DEM's or better techniques for detecting craters in optical images, as well as more complete labels, to detect craters across different size scales.
- DeepMoon is somewhat sensitive to domain shift, with decreased performance when evaluated on Mars DEM's. Despite the different properties of Mercury and Mars, however, DeepMoon can detect many craters.