DETECTING MOONQUAKES USING CONVOLUTIONAL NEURAL NETWORKS. F. Civilini[1] and R.C. Weber[1], [1]NASA Marshall Space Flight Center, 320 Sparkman Drive, Huntsville, AL (francesco.civilini@nasa.gov)

Introduction: Research in planetary seismology is fundamentally constrained by a lack of data. Seismological science products of future missions can typically only be informed by theoretical signal/noise characteristics of the environment [1] or likely Earth-analogues [2]. Although objectives can be re-assessed after some initial data-collection upon lander arrival, transfer of high-resolution data back to Earth is costly on lander power usage.

Over the last several years, development of GPU computing techniques and open-source high-level APIs have led to rapid advances in deep learning within the fields of computer vision, natural language processing, and collaborative filtering. These techniques are actively being adapted for a variety of tasks in seismology, including: earthquake detection [3], seismic phase discrimination [4], and ground-motion prediction [5].

Until the recent detection of marsquakes during the Mars InSight mission, the only other measurements of seismicity recorded outside of Earth was on the Moon during the Apollo missions (1969 to 1977) [6]. This dataset has been periodically revisited using new seismological methods, including ambient noise interferometry [7] and Hidden Markov Models [8].

We developed a binary seismic detection classifier using Convolutional Neural Networks (CNNs) trained from Earth seismic data and tested it against cataloged moonquakes recorded by the Apollo Passive Seismic Experiment (PSE). Two- to five-layer convolution models were tested against a subset of 200 Grade-A events from the PSE and obtained station accuracy average of 89-96%. The three-layer model was then used to catalog moonquakes from the Apollo 17 Lunar Seismic Profiling Experiment (LSPE). The algorithm was able to obtain detections for LSPE moonquakes with approximately one order of magnitude greater accuracy than a recent study using Hidden Markov Models (HMMs) [9].

Methodology: We built a prototype deep learning classifier that was able to distinguish between seismic activity and noise through examples of spectrogram images for each category recorded on Earth.

Earth seismic data was downloaded using the *IRIS* utility *PyWeed* in a time window around the earthquake first arrival. For the classifier, we used an interval of 180 seconds before and 20 seconds after the P-wave arrival of earthquakes greater than M_w 3 from the Piñon Flats Observatory (PFO) seismic station [Fig. 1]. This station was chosen due to its long operating duration and location between the San Andreas and San Jacinto fault zones. The variety of recorded events had diverse spectrogram characteristics which promoted algorithm generalization.

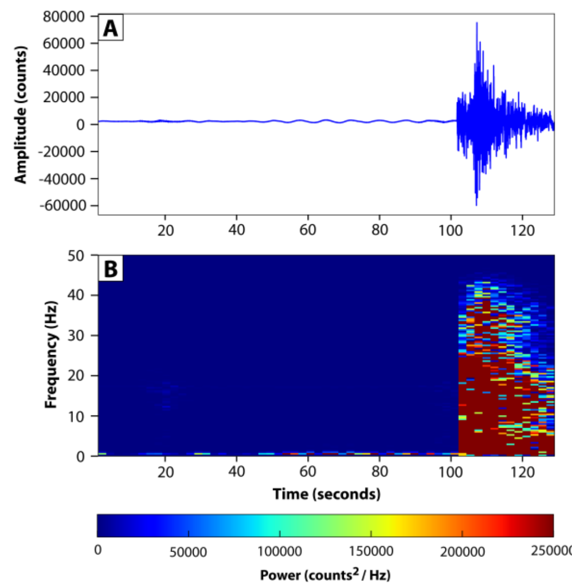

Figure 1: Time series [A] and spectrogram [B] of a M_w 3.7 earthquake at PFO.

Data augmentation is a technique in computer vision where new images are created by modifying existing data [10]. In image recognition, this is typically done by cropping, zooming, or rotating images. However, applying data augmentation in this manner will decrease the accuracy of our model, as cropped spectrograms may omit valuable information in the low or high frequencies. Instead, we chose twenty sliding windows across the noise and earthquake segments with one second overlap starting at 0 seconds for the noise and at 81 seconds for the earthquake (19 seconds prior to the onset of the P-wave at 101 seconds). A total of 27,800 spectrograms were used in the prototype, approximately 20% of which (5240) were separated into a validation set.

The prototype was built using the *fastai* computer vision library [10] with a batch size of 64 and image reduction to 224x224 pixels. Two- to five-layer convolution models with dropout were trained over fifty training cycles [Figure 2]. Each of the four models had an accuracy of over 99.9% on the validation set. We chose to name our ensemble of models "MoonNet".

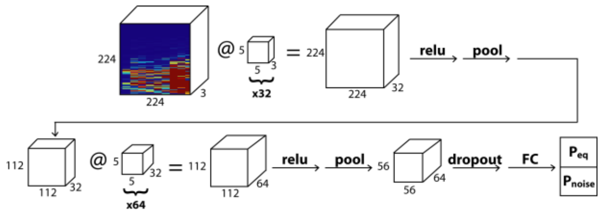

Figure 2: Architecture for the 2-layer CNN used in this study. An input spectrogram of size 224x224 with depth 3 (RGB) is convoluted with 32 kernels of dimensions 5x5x3 then passed through a Rectified Linear Unit (ReLU) activation function before undergoing dimensionality reduction (pooling). The resulting activation maps are then passed through this process again, dropout is applied, and a fully-connected layer (FC) is used to sort the volume into a two-element vector of quake or noise probabilities (Peq and P_{noise}).

Results: Each of the models was tested against 200 "Grade-A" events of the PSE and found to have between 89-99% accuracy across all seismic stations. An algorithm was developed to sort detections from each model into arrivals which removed false picks and extra consecutive detections. The three-layer convolution model was found to have the highest percentage of correct detection and was applied to the LSPE dataset.

This model was applied to approximately 8 months of LSPE data and compared with a recent study using HMMs [9] [Figure 3]. Two pre-processing steps were necessary to account for differences between the Earth-trained model and the lunar data: (1) the typical moonquake signal is much weaker in power compared to earthquakes and had to be capped at 1e-6 counts^2/Hz, and (2) a 10 Hz highpass filter was applied to remove low-frequency lunar noise. Figure 3 displays a typical one-hour trace of data and compares the HMM detections with MoonNet3L. We can observe through visual inspection that the results from the MoonNet3L model produce more accurate arrivals with fewer false or duplicate detections. The arrivals in each hour segment were summed across each day. The daily moonquakes vary with periodicity of approximately a month, consistent with the day/night cycle that drives the generation of thermal moonquakes. Similar to what was observed in the PSE application, there are occasional false detections in the LSPE catalog due to instrument glitches which register as large bursts of energy across the frequency spectrum. As hand-picked catalogs of the LSPE thermal moonquakes do not exist, it is difficult to quantify the accuracy of the model.

Through this study, we have demonstrated that accurate seismic detections of planetary seismicity can be made using deep learning generalized algorithms despite a lack of local training data. We hypothesize that such algorithms are lightweight enough to conduct lander-side event detection. Future work will assess the detection capability of these types of algorithms to other planetary bodies.

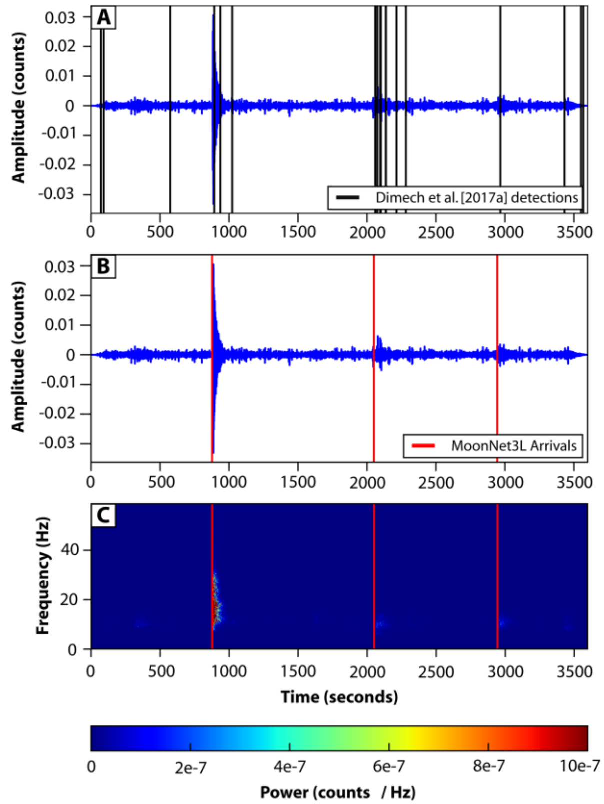

Figure 3: LSPE detection comparison between a Hidden Markov Models study [9] [A, black lines] and MoonNet3L [B, C, red lines].

References:

[1] Panning M. et al. (2018) *JGR: Planets, 123,* 167-179. [2] Zhan, Z. et al. (2014) GJI, *196*, 1796-1802. [3] Meier, M. et al. (2019) *JGR Solid Earth, 124*, 788-800. [4] Ross, Z. et al. (2019) *JGR Solid Earth, 124*, 856-869. [5] Trugman, D. and Shearer, P. (2018) *BSSA, 108*, 929-945. [6] Lognonné, P. (2005) *Annual Reviews of Earth Planet Science, 33*, 571-604. [7] Larose et al. (2005) GRL, 32 (16). [8] Knapmeyer-Endrun, B. and Hammer, C. (2015) *JGR: Planets*, *120*, 1620-1645. [9] Dimech, J. et al. (2017) *Results in Physics*, *7*, 4457-4458. [10] Howard, J. et al. (2018) fastai. [11] He, K. et al. (2016) *2016 IEEE Conference on Computer Vision and Pattern Recognition.*