PIXEL-WISE CLASSIFICATION AND AUTONOMOUS IMAGE ANALYSIS IN A REAL-TIME ROVER OPERATIONS SCENARIO: LESSONS LEARNED FROM THE CANMOON ANALOGUE MISSION.

A. D. Pascual[1,3], J. Kissi-Ameyaw[1,3], G. R. Osinski[2,3], K. McIsaac[1,3]. Department of Electrical and Computer Engineering, University of Western Ontario, [2]Department of Earth Sciences, University of Western Ontario , [3]Institute for Earth and Space Exploration, University of Western Ontario.

Introduction: The CanMoon lunar sample return analogue mission was a part of the Canadian Space Agency's (CSA) Lunar Exploration Analogue Deployment (LEAD) initiative, which aims to develop technologies and processes, as well as to train students and young professionals for future space missions. This mission was carried out by both the University of Western Ontario (Western), and the University of Winnipeg. The analogue mission was a simulation of a real-time operations scenario where a lunar-based scientific rover is operated and controlled from an Earth-based mission control center. For a full overview of the 2019 CanMoon analogue mission see Marion et al. [1]. Three teams were formulated for the mission: planning, science, and field teams. While the science team is in charge of deciding which scientific measurements to take and analyzing the returned data, the planning team commands the rover, keeping in mind the limitations imposed by the rover and the environment. Finally, the field team performs the activities of the rover in the exploration site.

As part of the mission, we explored the use and implementation of machine learning and deep learning models to aid in the decision making process of both the science and the planning teams. Leveraging the amount of imagery that the rover sends back to the ground, machine learning models could identify and localize the presence of different objects within the rover's surroundings. In addition, by performing the inferences within a short amount of time, these methods could provide quick insights and help guide the operators in making their decisions. It should be noted however, that the aim is not to replace the scientists and operators in making decisions, as this was a real-time mission with a human in the loop at all times.

Of the 4 science goals laid out by the science team [2], one was finding xenoliths to determine if a rover would be able to identify and sample pieces of lunar mantle material. Keeping in line with this goal, it quickly became apparent that this same task could potentially be accomplished by the models. Thus, the models were designed to perform pixel-wise classification on imagery from the rover. Even though the team had produced geological maps of the landing site based on satellite imagery [2], there were no images on the ground that would have allowed for pretraining a classifier prior to the mission. This also follows the scenario of a real mission where there will be no prior imagery of the ground from the rover's perspective.

It came as a surprise to the team when the first panorama image from the rover contained something that did not appear on the remote sensing data: lichen. Since lichen was protected at this field site, a limitation was imposed on the rover that it could not traverse across lichen. Here then was another opportunity for the use of autonomous image classification, i.e., finding where the lichen was within an image.

Methodology: Two algorithms were implemented for this task: random forests [3], and artificial neural networks [4]. To build the training data, all the images from the first few days where curated by the team. Zoom images and panorama images containing known xenoliths and lichen were used to label xenolith pixels and non-xenolith pixels. A sample panorama image is shown in Figure 1 and a sample zoom image is shown in Figure 2. Notice that xenoliths and lichen are visually very similar to one another. And where humans cannot differentiate between one and the other, machine learning models might be able to do better.

Figure 1. One of the images returned from a panorama prior to stitching.

Figure 2. Targeted zoom image of what appears to be a xenolith.

TextureCam. An existing implementation of pixel wise classification with random forests was used [3, 5]. The process involved three main steps:

1. Training image preparation - image filtering and remapping to HSV
2. Training a Random Forest classifier – an ensemble of decision trees
3. Inference on new image scenes.

The TextureCam software has been used before in the CanMars 2016 Mars Sample Return Analogue mission and has proved to be an effective classifier when distinguishing between different colored units within a scene [6]. One important distinction between the CanMars and the CanMoon mission is that CanMars was not a real-time mission. There was a significant time delay between uplink to the rover and the downlink of returned data.

Artificial Neural Networks. An Artificial Neural Network (ANN) is a classifier that takes in an input vector, performs a linear combination of the input, and passes it to a non-linear activation function [4]. This process is repeated for however many nodes one layer has. For each layer in the network, the outputs of the activation functions are linearly combined again to be fed to another non-linear activation function. The choice of how many nodes and how many layers are hyperparameters that is chosen beforehand. The network then outputs a probability for which class the output belongs to. Similar to the TextureCam implementation, raw pixel values in the HSV color space has been chosen as the input as this mapping has been known to perform well in classification tasks. The output layer then outputs probabilities for each pixel to belong to one of the following classes: xenolith, lichen, scoria, basalt. Labels for each class has been done manually from zoom images. The architecture composed of 2 layers, and 32 nodes for each layer. The sigmoid function has been chosen as the activation for each node, and Adam was chosen as the optimizer.

Results: The ANN performed quite well, achieving an 88% accuracy after being trained for 100 epochs. Figure 4 shows the predictions made on Figure 1. The resulting heat map shows that the classifier cannot pinpoint exactly where each object is going to be, but it can provide a rough map of where to expect lichen and xenoliths. Once used to highlight where the lichen could be, the algorithm proved to be helpful during traverses when lichen was illuminated by direct sunlight, making them difficult to see. Figure 4 shows a heat map obtained by the ANN.

Figure 5 shows the predictions made by TextureCam on Figure 2. Right away, it can be seen that the xenolith is highlighted in the heatmap, demonstrating the potential of the algorithm in finding objects of interest.

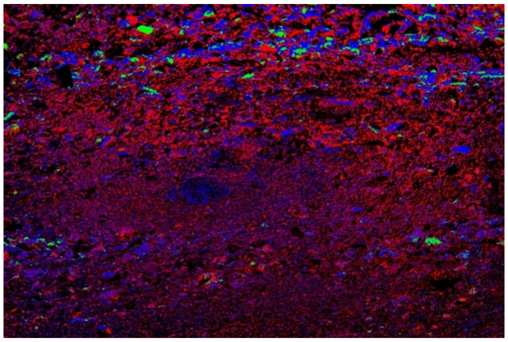

Figure 4. Predictions made by the ANN on Figure 1. Lichen were highlighted as the green pixel.

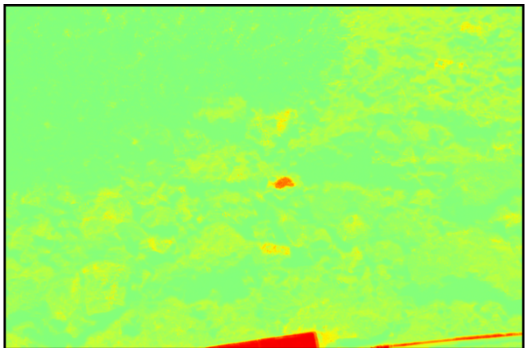

Figure 5. Predictions made by TextureCam on where the Xenolith is within the image. Xenoliths were highlighted as the red pixel.

Conclusions. There really is no computer more powerful than the human brain especially when it comes to finding objects within an image. Thus, in a real-time operations scenario, a machine learning model would not be able to outperform scientists in finding xenoliths accurately. However, in the same real-time scenario, copious amounts of imagery is being returned by the rover. Machine learning models could then swiftly make inferences on the images and help provide scientists with quick insights and draw their attention towards areas with objects of interest, maximizing the opportunity for science throughput.

References: [1] Marion C. et al. (2020) LPSC LI, this conference. [2] Morse, Z. et al. (2020) LPSC LI, this conference. [3] Breiman, L. Machine Learning (2001)45:5. https://doi.org/10.1023/A:1010933404324 [4] Daniel G.G. (2013) Artificial Neural Network. In: Runehov A.L.C., Oviedo L. (eds) Encyclopedia of Sciences and Religions. Springer, Dordrecht. [5] Thompson, D. R., et al. "TextureCam: Autonomous Image Analysis for Astrobiology Survey." Lunar and Planetary Science Conference. Vol. 43. 2012. [6] Kissi J. K. et al. (2017) LPSC XLVIII