Enabling Autonomy and Operations for Lunar Surface Missions: An Overview of Demonstrated Capabilities

Matt Cross[1] and Andrew. J. Macdonald[2], B. Bonham-Carter[3], H. Burd[3], L. Chavier[3], T. Heydrich[3], S. Pillay[3], G. O'Shea[3], E. Smal[3], M. Maharib[3], K. Raimalwala[3], M. Faragalli[3], M. Battler[3]

[1]*Mission Control, 162 Elm St W, Ottawa, K1R 6N5 Canada, cross@missioncontrolspaceservices.com*

[2]*Mission Control, 162 Elm St W, Ottawa, K1R 6N5 Canada, macdonald@missioncontrolspaceservices.com*

[3]*Mission Control, 162 Elm St W, Ottawa, K1R 6N5 Canada*

ABSTRACT

Early lunar micro-rover missions will have constrained mission life and downlink capacity, which limits the data volume that is received by ground operators for operational decision making. Increasing operational autonomy is required to maximize the scientific, operational, and economic return to ensure the viability and sustainability of lunar and planetary surface missions. To enable operational autonomy, Mission Control has developed MoonNet, an AI-enabled terrain classifier, to perform image segmentation on low-power spaceflight computing. The outputs of the image segmentation can be used to support data compression, downlink prioritization, and automated instrument targeting. This paper presents the development and deployment of MoonNet for demonstration as part of the ispace M1 mission.

1. INTRODUCTION

1.1. Challenges in Commercial Lunar Rover Missions

Early commercial lunar surface missions will be limited to one lunar day (14 Earth days), due to the harsh environment, and nominal operations at mid/low latitudes will likely be 10-12 Earth days. Lander payloads, including rovers that relay upon the lander for communications relay, must share a constrained downlink capacity. For rovers that relay upon direct telecommanding, consistent and frequent data is required for operational decision making. Delays in receiving data result in less operational delays and limit the scientific, explorative, and economic outcomes of the rover mission. These constraints motivate the need for innovative concept of operations and technologies to ensure viable rover missions.

1.2. Mobility and Science Operations

Several factors necessitate increased autonomy in mobile science operations. Early Mars rover operations relied upon ground-in-the-loop visual surface characterization and subsequent analysis for decision-making over one or more tactical cycles [1]. Recent advancements continue to push operational autonomy to allow for some operations to be conducted without ground-in-the-loop [2]. Upcoming Lunar rover missions, however, will have reduced latency, shorter lifetimes, along with constrained bandwidth. These factors necessitate rapid operational decision-making processes with limited data with limited time to analyze data, identify features of interest, and make science-driven decisions.

The NASA VIPER rover that will fly to the south polar region is a large rover (~300kg) but will have a constrained direct-to-Earth communications channel of 230 kbps [3]. Small-scale commercial Lunar rovers will also be constrained; payloads, including micro rovers on Astrobotic's Mission One will be allocated 10 kbps per kg according to standard payload data rate allocation advertised in their Payload User Guide (PUG) [4]. As per their CubeRover PUG, a 6kg payload will be allocated 60 kbps [5].

Sensors are growing increasingly powerful; however, data transfer rates are not yet sufficiently high to downlink high volumes of data for near-term operational decision making. Intelligent methods for selection, compression, and prioritization of science data collected by the rover are needed to maximize scientific, explorative, and economic return, and to facilitate operational decision making.

The nature of scientific discovery makes onboard autonomy compelling. It increases the chances of detecting valuable novel/sparse features that may otherwise be missed in scenarios that prioritize driving and other mission needs. For example, NASA's Opportunity rover drove 600 ft past the Block Island meteorite, one of its biggest discoveries, before the science team discovered it and decided to drive back to investigate it [6].

Rover missions will benefit from autonomy in data processing and decision making with short duration tactical operational cycles and pressure to achieve science objectives. MoonNet, an AI-enabled terrain classifier developed by Mission Control, advances capabilities for lunar applications from its ASAS-CRATERS (Autonomous Soil Assessment System: Contextualizing Rocks, Anomalies and Terrains in Exploratory Robotic Science) [7] field study. MoonNet enables such capabilities with the goal of maximizing scientific return in upcoming missions.

1.3. State-of-the-Art in Autonomous Perception for Planetary Science

The state-of-the-art in terrain classification leverages high performance Convolutional Neural Networks (CNNs) that find natural features and complex patterns in the image. Soil Property and Object Classification (SPOC) [8] has a terrain classifier that uses Fully CNNs (FCNNs). SPOC was further extended though Machine learning-based Analytics for Automated Rover Systems (MAARS), which included automated scientific captioning of terrain images [9]. Gonzalez and Iagnemma [10] published a comparative analysis of CNNs, Deep Neural Networks, and classical algorithms such as Support Vector Machines (SVM). These and other works have focused on classifying Mars surface images to improve autonomy for Mars rovers. Chiodini [11] evaluated the CNN Deeplab-v3+ for labeling a voxelized stereo pair to differentiate rocks from sand on a semantic model of the scene in front of the rover.

For Lunar applications, terrain classification motivated by scientific research has focused on crater detection using orbital data. Stepinski et al. and Chung et al. offer a review of traditional machine learning techniques, including SVMs [12], [13]. Silburt et al. [14] explored the use of CNNs to detect craters using a digital elevation model merged from *Lunar Reconnaissance Orbiter* and *Kaguya* data. For the lunar surface, Matthies [15] investigated using CNNs for crater detection with monocular camera imagery as input to a crater-based rover localization.

Kerner et al. [16], [17] demonstrated the capability to detect novel geological features in multispectral images of the Martian surface. They show that a spatial-spectral error map can enable both accurate classification of novelty in multispectral images as well as human-comprehensible explanations of the detection. Stefanuk and Skonieczny [18] demonstrated using variational autoencoders for novelty detection in planetary datasets.

These studies have successfully demonstrated the use of deep learning to improve terrain classification of images from Mars rover datasets or from a laboratory setting. Recent work by Mission Control holistically studied terrain classification in a real-time system for a science-driven rover mission and its implications on mission operations [19]. The Mission Control terrain classifier was first developed under the CSA-funded Autonomous Soil Assessment System (ASAS) [20]. In 2019, it was used onboard a rover to classify eight Mars-relevant terrain types in real-time at ~15 FPS as the rover drove at 20cm/s at a high-fidelity analogue site in Iceland, with a sample result shown in Figure 1. This study was a part of SAND-E (Semi-Autonomous Navigation for Detrital Environments), a NASA PSTAR (Planetary Science and Technology Through Analog Research) funded project to inform Mars2020 operations [19].

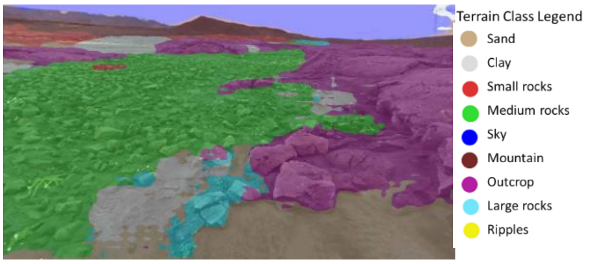

Figure 1: Result from field-testing the deep-learning based terrain classifier in Iceland.

ASAS-CRATERS further advanced ASAS to classify terrain into pre-determined geological categories and detect novel geologic features [7]. ASAS-CRATERS was deployed onto a Xiphos Q8 computing hardware, which has a similar low-power and low-mass profile as the spaceflight-qualified Q7S.

Modern AI, enabled by recent advances in deep learning, has the potential to help space operations to become more autonomous, adapt to unexpected changes, synthesize data for human analysis, and maximize productivity during idle time. Mission Control, as part of a NASA-funded PSTAR project, has already shown the value of deep learning in an analogue planetary exploration mission [19]. Section 2 of this paper presents the development and early results of MoonNet leading up to its planned lunar demonstration.

1.4. Overview of Lunar Capability Demonstration

MoonNet was manifested as a payload onboard the first ispace mission M1. Also manifested on M1 was the Emirates Lunar Mission (ELM) *Rashid* rover led by the Mohammed Bin Rashid Space Centre (MBRSC) [21] [22]. MoonNet was intended to classify lunar surface features visible in images from *Rashid* as part of a collaboration with ELM. The key objectives were to demonstrate MoonNet in flight during the surface operations of ELM and to demonstrate an increase in performance following in-flight model weights update. These objectives were to be achieved by receiving images from ELM, ingesting these images into Mission Control's deep learning model training pipeline [23] and ultimately deploying a retrained version of the MoonNet model to the lunar surface retraining the MoonNet model.

The ELM *Rashid* rover had two identical navigation cameras: CAM-1, a forward-facing navigation camera mounted on a mast gimbal, and CAM-2, a rearward-facing fixed-mounted camera [24]. The received CAM-1 and CAM-2 images were to be used for retraining MoonNet on the ground prior updating the model weights.

MoonNet development began with user needs discussion and labelling campaign in March 2022. Model development, including augmentation strategy and training, was finalized and tested during assembly, integration and testing (AIT) in August 2022. MoonNet was then launched on the M1 spacecraft on December 11th, 2022. After M1 entered lunar orbit in late-March 2023, the MoonNet payload was confirmed to be operating nominally making it the world's first deployed deep learning AI in lunar orbit.

MoonNet was to be the world's first demonstration of deep learning on the lunar surface, a historic milestone for space exploration. However, the M1 spacecraft, did not successfully soft-land on the lunar surface [22]. Despite the unsuccessful landing, there were achievements and valuable lessons for operational autonomy in the development of MoonNet software and its concepts of operation, which are described in the following section.

2. OVERVIEW OF MOONNET DEVELOPMENT

The ASAS-CRATERS program concluded at the end of 2021, at which point the knowledge developed over years of working on the technology was applied to the task of deployment on spaceflight computing hardware for demonstration in lunar surface operations. Modification of the ASAS-CRATERS models and down selecting a single, trained model to become the deployed flight model, MoonNet, required thinking about the specific requirements of the mission.

2.1. Science User Needs

ASAS-CRATERS provided the groundwork user needs analysis, with refinement required for the details of the particulars of M1. AI Specialists, software engineers, and planetary scientists established the user needs of the demonstration following best practices for Machine Learning model development [25], which resulted in a labelling taxonomy to train MoonNet. With input from planetary geologists, AI Specialists downselected seven class labels for the taxonomy based on the most successful classes for the various models that had been tested and an understanding of the data available: crater interior, crater rim, regolith, rocks, rover tracks, sky, and spacecraft.

2.2. Labelling Campaign

One of the significant challenges of this demonstration was the lack of representative imagery: there were no images from the M1 landing site surface at the height of the *Rashid* camera with the specific camera sensor available for labelling and training. The difference in optics, sensor spectral range, lighting conditions, and geologic conditions can impact the performance of the neural network. To overcome these challenges, a two-stage training strategy was adopted starting with images from Mission Control's Moonyard collected for ASAS-CRATERS and then finetuning on lunar imagery from the Chang'E missions [26]. Both datasets were labelled using the MoonNet taxonomy on the Labelbox platform. Representative examples from Labelbox are shown in the following Figure 2:

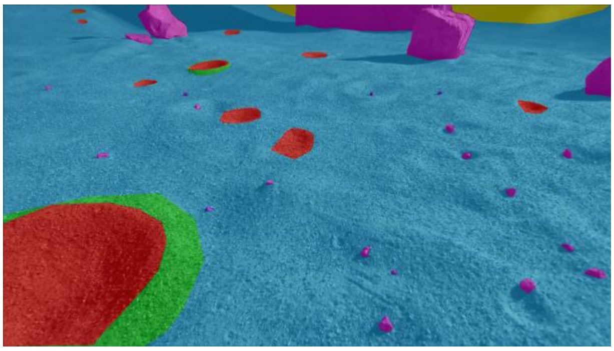

Figure 2: Examples of a labelled image from Mission Control's Moonyard on the Labelbox platform.

Mission Control AI Specialists performed quality control and assurance on the labelled data. There were consistent labels applied across rocks, sky, and spacecraft. In areas of flat lighting conditions some craters or rocks were missed and corrected when possible. Crater rims were not clearly distinguishable compared to the Moonyard. In several instances, craters received the rock label, this was caught on review. The labelling campaign resulted in 1519 Moonyard images and 275 Chang'E images.

2.3. Training

A custom training regimen with specific augmentations was developed to increase MoonNet's performance during lunar operations. This included custom aspect ratio cropping, image flipping, and colour jitter to match the expected concept of operations and properties of *Rashid* CAM-1. Custom data products and metrics were used to assess performance, including mean Intersection-

Over-Union, class-based accuracy, and overlays like those shown in Figure 3.

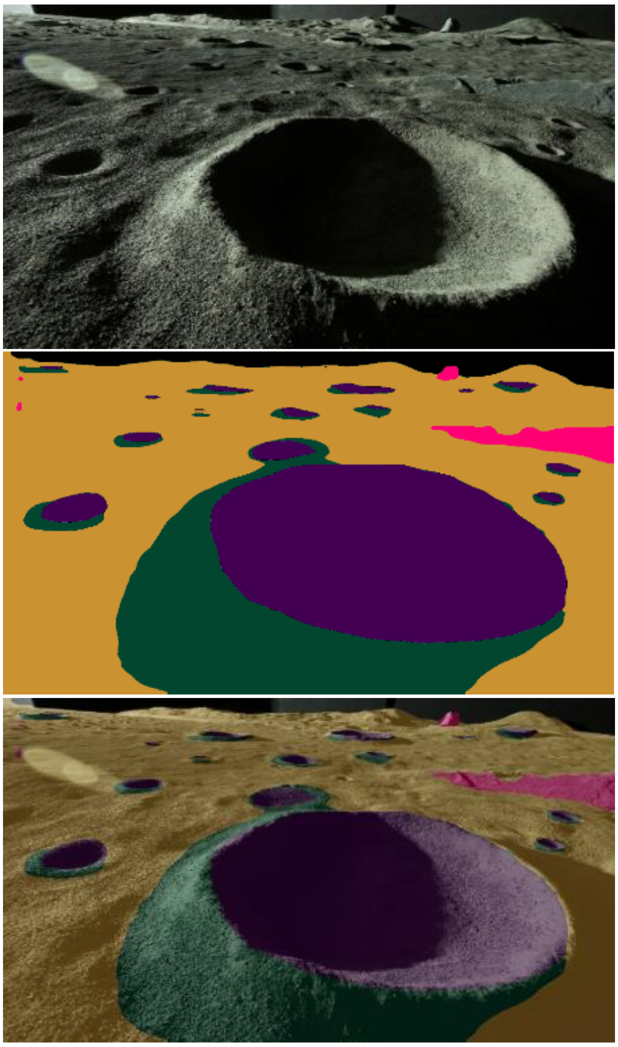

Figure 3: Input image to MoonNet from the Moonyard (top). Output MoonNet prediction (centre) overlaid on original image (bottom).

2.4. Deployment to and Testing on Spaceflight Computing Hardware

MoonNet was deployed onto a Xiphos Q7S as a payload on M1. The Q7S to listen to the data stream for images from CAM-1 for processing with MoonNet. The MoonNet software ran entirely on the Q7S ARM Cortex-A9 600 MHz CPU. The Q7S has 512 MB RAM and a power draw of approximately 1 W. The limited computational power and memory was another challenge of this demonstration, as CNNs are not ordinarily run on low-SWaP hardware. The radiation environment on route to the lunar surface and during surface operations posed an additional challenge [27].

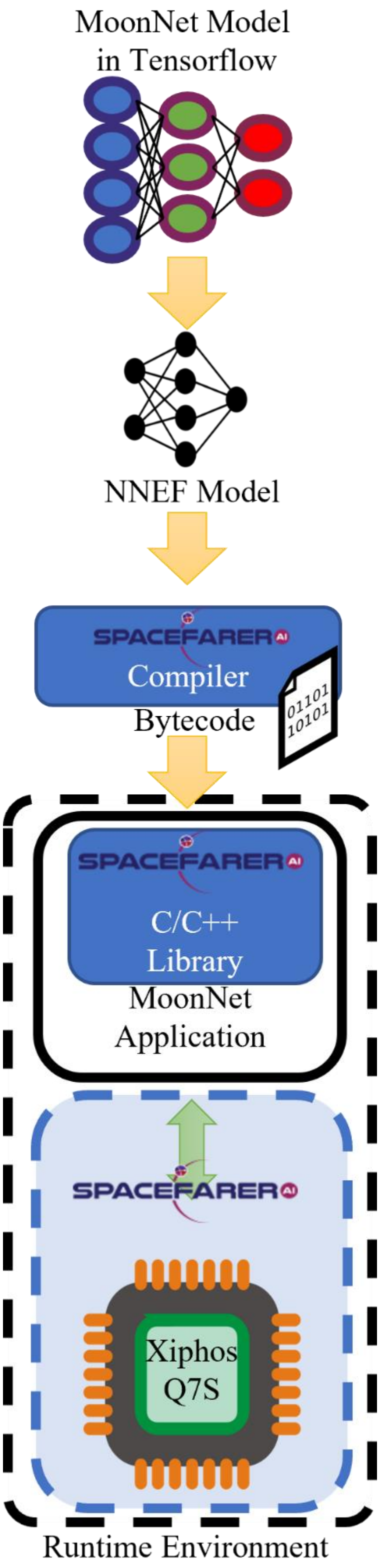

Figure 4: The Spacefarer AI Deployment Toolkit accelerates high-level trained models into production of edge devices for space flight applications, moving from training to compiling to running the model.

2.5. Ground Segment and MLOps Development

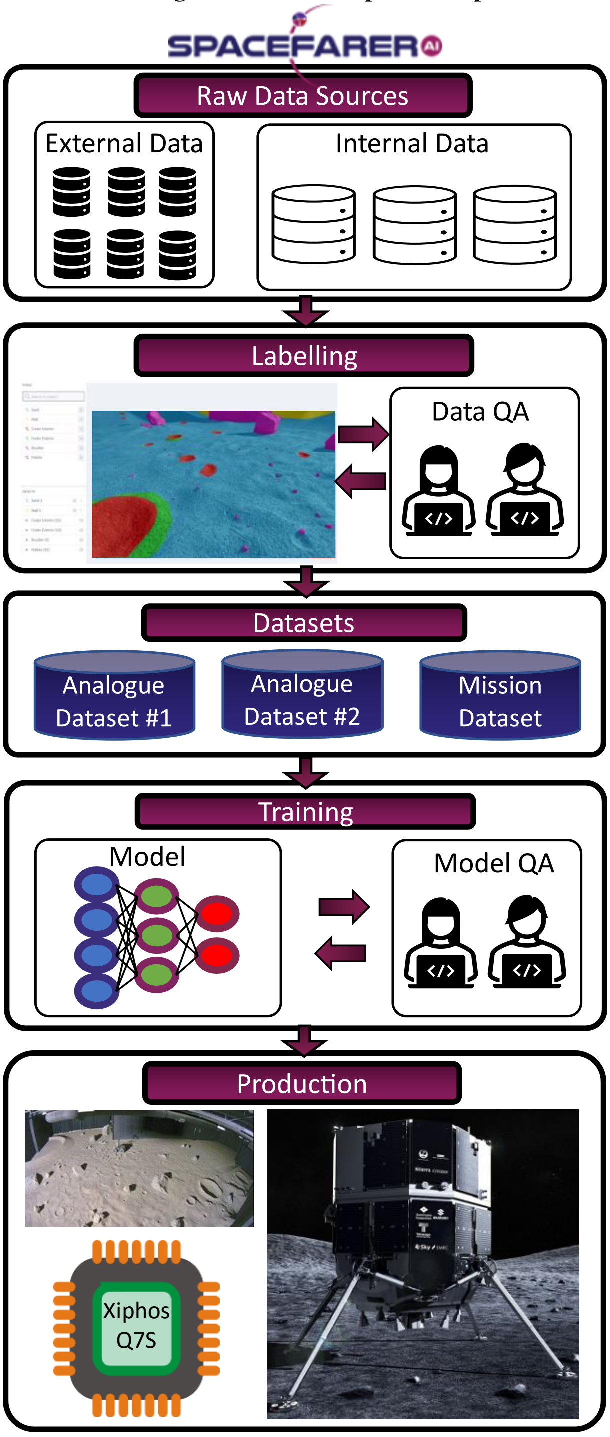

Figure 5:High-level view of Mission Control's deep learning pipeline used to develop and deploy MoonNet onboard a Q7S. Image of ispace HAKUTO-R lander credit: ispace.

With the challenges of *a priori* training data and computational limitations, it was known that MoonNet as first deployed would have limited performance. The payload software was developed to permit updates to the MoonNet model weights to improve the performance based on real data acquired during surface operations. There were two opportunities to stage and uplink an update of the model weights. The model update strategy, a distinguishing feature of MLOps, required Mission Control to build the first ground infrastructure and workflow required to update an onboard neural network during lunar operations. This included the construction of a deep learning pipeline for processing, verifying, and training on new data [23], as shown in Figure 5.
Mission Control developed custom software and pipelines to support real-time evaluation and assessment of MoonNet performance during lunar surface operations. This included detailed experimental assessments between different model training runs to improve MoonNet performance on the flight model during the 10-12 days of surface operations through the deployment and staging of new models trained on CAM-1 data.

Mission Control built an automated validation and verification pipeline that ensured proper function and performance of the newly trained MoonNet models by running comparisons between the high-level training algorithms and performance on the in-house Engineering Model. This pipeline allowed a robust, reproducible staging pipeline for sending verified models to ispace during uplink windows.

Surface operations were planned such that MoonNet model would be updated in flight twice to demonstrate improved model performance during flight, demonstrating the importance of a robust MLOps pipeline for production level edge AI applications in space.

2.6. Results

Mission Control achieved the following Milestones as part of this capability demonstration:

- MoonNet model training and deployment of AI model as flight software on flight hardware.
- Successful AIT of MoonNet with ispace HAKUTO-R lander and MBRSC *Rashid* rover flight models
- Passing final software tests with Mission Control payload inside the SpaceX rocket fairing in Florida.
- Launch on SpaceX Falcon9.
- Establishing nominal payload operation post-launch.
- Nominal operations of payload through three months of deep space travel, including correct recovery from radiation strikes.
- Contact with payload in lunar orbit, confirming nominal operation.
- Successful final rehearsal of machine learning lunar operations pipeline with Spacefarer AI.

3. USE CASES FOR MOONNET IN LUNAR ROVER OPERATIONS

MoonNet itself is an enabling tool for enhancing operations and operational autonomy of rover missions. This section presents use cases, both demonstrated and conceptual, for MoonNet in lunar rover operations.

3.1. Feature-Based Image Compression

Figure 6 shows a demonstrated example of one such product derived from the MoonNet outputs. Fig. (a) shows an input image from the Moonyard containing sand, craters, rocks, and the background wall. Science users may have mission objectives tied specifically to craters and rocks on the lunar surface and so the segmentation map prediction is used to generate a smaller image product that keeps the high-resolution of these features while down sampling the background wall and foreground sand, as shown in the output image (b). In future missions these kinds of derived data products can be constructed onboard and used to maximize the science return in limited communications bandwidth scenarios.

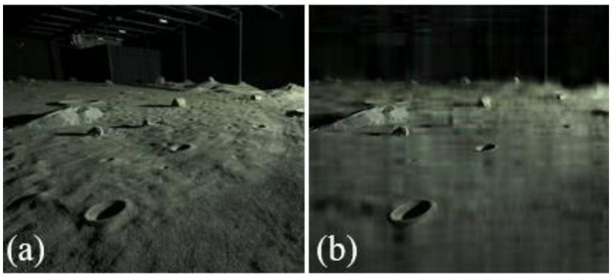

Figure 6: An example of a derived data product from the MoonNet predictions from an input image (a). The output prediction is used to selectively down sample certain image features to create a smaller image while maintaining high fidelity for craters and rocks (b).

3.2. Downlink Prioritization

An alternative concept is to the segmentation-based image compression to assess the features contained within an image and to prioritize the downlink of images based on desired features. The rover system can collect high-quality images at a higher rate than it can downlink to ground. Science users may have mission objectives that prioritize craters on the lunar surface. Images that contain these features can be prioritized for downlinking over images that do not. An example is shown in Figure 7.

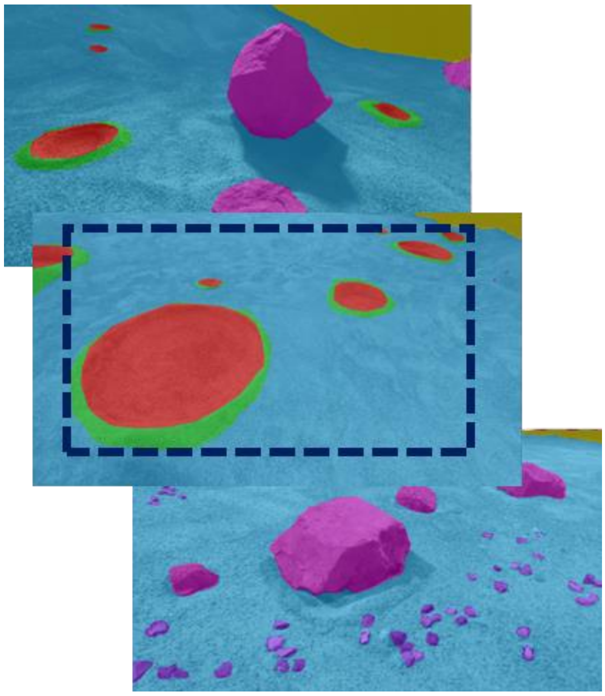

Figure 7: An example concept application for derived data products. The output prediction is used to select and prioritize images that contain desirable features.

3.3. Supporting Instrument Targeting

MoonNet segmentation outputs can support science users in targeting science features of interest for rover missions that have additional science instruments that can conduct targeted investigations of features, such spectrometers. The science user on the ground is presented with the MoonNet predictions overlayed onto the base image and have the feature classes labeled in a user interface, as exampled in Figure 8. The science user can then quickly select and communicate the selection to the rover operators to perform the targeted investigation.

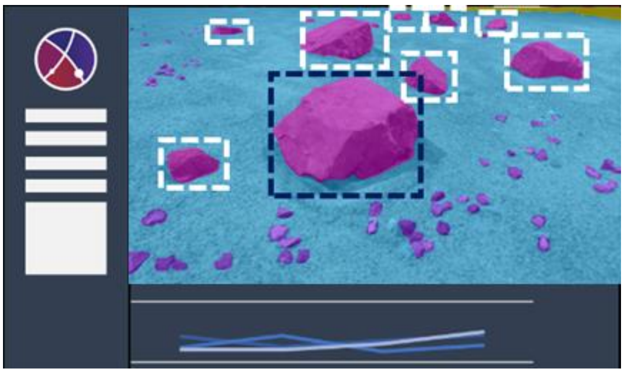

Figure 8: A mock-up of a conceptual user interface that can support processing autonomously identified features for tasks such as instrument target requests.

3.4. Enabling Autonomous Targeting

The advancement of user targeting is to enable onboard autonomous instrument targeting. In this use case, the MoonNet predictions can be used as an input directly onboard the rover to identify a target of interest as pre-defined by science user needs. The estimated relative position of the feature of interest, for example a rock, can then translate to mechanical actuation to point the instrument, for example a spectrometer, at the science feature, and once in position the instrument can acquire its measurements, such as with AEGIS (Autonomous Exploration for Gathering Increased Science) on Mars Science Laboratory [29], as exampled in Figure 9. A further advancement is to combine this autonomous targeting with downlink prioritization, where the highest value science data products are prioritized for downlink to the science user.

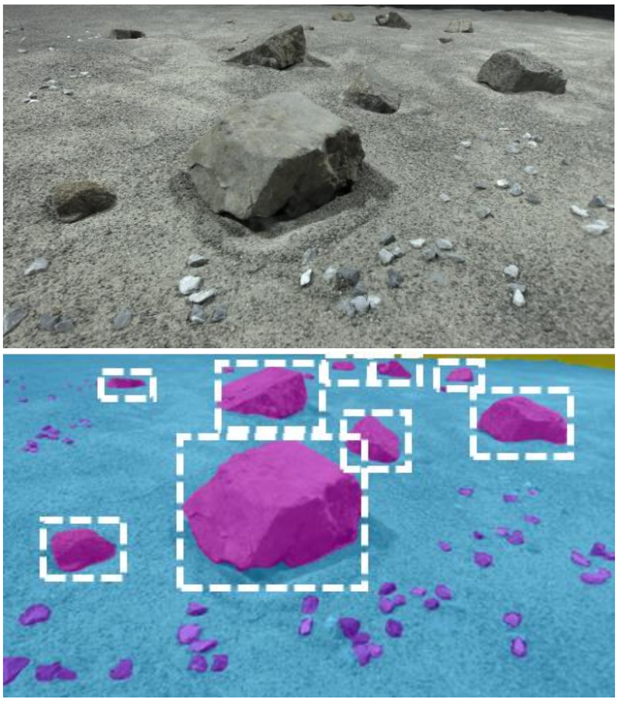

Figure 9: An example concept application for the derived data product. The output prediction is used to target instruments onto features of interest to acquire the science data without requiring the science user on the ground in the loop.

3.5. Terrain Classification Mapping Aggregation

An initial objective of ASAS was to assess soil properties for mobility and aggregate those soil assessments onto a map generated by the rover. This concept is further extended to aggregate the segmented terrain features and science targets onto a mapping product.

4. DISCUSSION AND CONCLUSSIONS

Mission Control developed demonstrate capabilities with an in-flight AI model. Future work will be to deploy advanced AI models directly onto a rover system. This direct deployment with a rover system will allow Mission Control to advance operational autonomy capabilities to overcome operational constraints to maximize the scientific, operational, and economic return of lunar rover missions. Additionally, the AI deployment capabilities developed for this demonstration are being applied to a variety of domains including Earth observation, space domain awareness, and AI-enabled robotic arms.

The development of MoonNet for deployment on low-power was a major achievement towards enabling CNN-based autonomy applications in deep space missions.

5. ACKNOWLEDGEMENTS

The authors thank and acknowledge the support of the Canadian Space Agency [3CAPDEMO21], and the demonstration partners for science and technology support: Dr. Ed Cloutis at the University of Winnipeg, Dr. Krzysztof Skonieczny at Concordia University. This work is part of the framework of the Emirates Lunar Mission (ELM). The ELM is conducted by the Mohammed Bin Rashid Space Center (MBRSC), in Dubai, UAE. The ELM is funded by the ICT Fund as part of the Mars 2117 initiative.

6. REFERENCES

[1] D. Gaines *et al.*, 'Productivity challenges for mars rover operations', in *Proceedings of 4th Workshop on Planning and Robotics (PlanRob)*, London, UK, 2016, pp. 115–125.

[2] V. Verma *et al.*, 'Autonomous robotics is driving Perseverance rover's progress on Mars', *Sci. Robot.*, vol. 8, no. 80, p. eadi3099, Jul. 2023, doi: 10.1126/scirobotics.adi3099.

[3] A. Colaprete, 'An Overview of the Volatiles Investigating Polar Exploration Rover (VIPER)', presented at the Missions and Instruments Virtual Workshop, SSERVI, Feb. 07, 2020. [Online]. Available: https://youtu.be/mUnkY_Rwo44

[4] 'Astrobotic - Payload User Guide.pdf'. Aug. 2019.

[5] 'Astrobotic CubeRover - Payload User Guide'. Apr. 2020.

[6] 'NASA - Meteorite Found on Mars Yields Clues About Planet's Past'. https://www.nasa.gov/mission_pages/mer/news/mer-20090810.html (accessed Sep. 18, 2020).

[7] K. Raimalwala, M. Faragalli, E. Reid, E. Smal, and M. Battler, 'Autonomous Soil Assessment

System: Contextualizing Rocks, Anomalies and Terrains in Exploratory Robotic Science (ASAS-CRATERS)', in *Lunar Surface Science Workshop 2020*, Denver, CO: Lunar & Planetary Institute, Apr. 2020.

[8] B. Rothrock, R. Kennedy, C. Cunningham, J. Papon, M. Heverly, and M. Ono, 'SPOC: Deep Learning-based Terrain Classification for Mars Rover Missions', in *AIAA SPACE 2016*, American Institute of Aeronautics and Astronautics, Sep. 2016. doi: 10.2514/6.2016-5539.

[9] M. Ono *et al.*, 'MAARS: Machine learning-based Analytics for Automated Rover Systems', in *2020 IEEE Aerospace Conference*, Mar. 2020, pp. 1–17. doi: 10.1109/AERO47225.2020.9172271.

[10] R. Gonzalez and K. Iagnemma, 'DeepTerramechanics: Terrain Classification and Slip Estimation for Ground Robots via Deep Learning', *ArXiv180607379 Cs*, Jun. 2018, Accessed: Aug. 17, 2018. [Online]. Available: http://arxiv.org/abs/1806.07379

[11] S. Chiodini, L. Torresin, M. Pertile, and S. Debei, 'Evaluation of 3D CNN Semantic Mapping for Rover Navigation', in *2020 IEEE 7th International Workshop on Metrology for AeroSpace (MetroAeroSpace)*, Jun. 2020, pp. 32–36. doi: 10.1109/MetroAeroSpace48742.2020.9160157.

[12] T. F. Stepinski, W. Ding, and R. Vilalta, 'Detecting impact craters in planetary images using machine learning', in *Intelligent data analysis for real-life applications: theory and practice*, IGI Global, 2012, pp. 146–159.

[13] C. P. Chung, C. O. Chung, and C. H. Yam, 'Lunar Crater Identification from Machine Learning Perspective', *Acta Futura*, p. 8, 2014.

[14] A. Silburt *et al.*, 'Lunar Crater Identification via Deep Learning', *Icarus*, vol. 317, pp. 27–38, Jan. 2019, doi: 10.1016/j.icarus.2018.06.022.

[15] L. Matthies *et al.*, 'Lunar Rover Localization Using Craters as Landmarks', in *2022 IEEE Aerospace Conference (AERO)*, Mar. 2022, pp. 1–17. doi: 10.1109/AERO53065.2022.9843714.

[16] H. R. Kerner, D. F. Wellington, K. L. Wagstaff, J. F. Bell, C. Kwan, and H. Ben Amor, 'Novelty Detection for Multispectral Images with Application to Planetary Exploration', *Proc. AAAI Conf. Artif. Intell.*, vol. 33, pp. 9484–9491, Jul. 2019, doi: 10.1609/aaai.v33i01.33019484.

[17] H. R. Kerner *et al.*, 'Comparison of novelty detection methods for multispectral images in rover-based planetary exploration missions', *Data Min. Knowl. Discov.*, Jun. 2020, doi: 10.1007/s10618-020-00697-6.

[18] B. Stefanuk and K. Skonieczny, 'Novelty detection in rover-based planetary surface images using autoencoders', *Front. Robot. AI*, vol. 9, 2022, Accessed: Sep. 06, 2023. [Online]. Available: https://www.frontiersin.org/articles/10.3389/frobt.2022.974397

[19] R. C. Ewing *et al.*, 'Overview and Initial Results of SAND-E: Semi-Autonomous Navigation for Detrital Environments', in *51st Lunar and Planetary Science Conference*, 2020.

[20] M. Faragalli, K. Skonieczny, and K. Raimalwala, 'Predicting Non-geometric Hazards Using Machine Learning for Safer Rover Mobility Operations', in *International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS) 2018*, Jun. 2018.

[21] S. Almaeeni, S. Els, and H. Almarzooqi, 'The Rashid rover: to guide the way for the next generation lunar missions and solar system exploration', pp. EGU21-14732, Apr. 2021, doi: 10.5194/egusphere-egu21-14732.

[22] G. Conroy and M. Naddaf, 'Private ispace Moon landing fails: researchers are investigating', *Nature*, vol. 617, no. 7959, pp. 18–19, 2023.

[23] A. J. Macdonald *et al.*, 'Enabling Autonomy with a Deep Learning Framework for Planetary Exploration', in *ASCEND 2022*, American Institute of Aeronautics and Astronautics. doi: 10.2514/6.2022-4295.

[24] Z. Ioannou *et al.*, 'The Image Acquisition Systems On-Board the Rashid Rover of the Emirates Lunar Mission', vol. 44, p. 285, Jul. 2022.

[25] People + AI Research Team, *People + AI Guidebook*. Google, 2021. [Online]. Available: https://pair.withgoogle.com/guidebook/

[26] China National Space Administration, 'Ground Research and Application System of China's Lunar and Planetary Exploration Program, Chang'E 4 Panoramic Cameras Dataset', 2020. https://moon.bao.ac.cn

[27] K. Raimalwala *et al.*, 'Enabling Autonomy in Commercial-Class Lunar Missions', in *International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS) 2020*, Oct. 2020.

[28] L. Chavier *et al.*, 'Deploying Artificial Intelligence Capabilities by Hybridizing a Neural Network on a Satellite', in *ASCEND 2023*, American Institute of Aeronautics and Astronautics.

[29] R. Francis *et al.*, 'AEGIS autonomous targeting for ChemCam on Mars Science Laboratory: Deployment and results of initial science team use', *Sci. Robot.*, vol. 2, no. 7, p. eaan4582, Jun. 2017, doi: 10.1126/scirobotics.aan4582.