

NOVELTY-DRIVEN ONBOARD TARGETING FOR MARS ROVERS

Virtual Conference 19-23 October 2020

Kiri L. Wagstaff¹ (kiri.l.wagstaff@jpl.nasa.gov), Raymond Francis¹, Hannah Kerner², Steven Lu¹, Favour Nerrise², James F. Bell III³, Gary Doran¹, and Umaa Rebbapragada¹

¹Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA, 91109, USA

²University of Maryland, College Park, MD, 20742, USA

³Arizona State University, Tempe, AZ, 85381, USA

ABSTRACT

Current Mars surface exploration is primarily prescribed on a day-by-day basis. Mars rovers have a limited ability to autonomously select targets for follow-up study that match pre-defined target signatures. However, when exploring new environments, we are also interested in observations that differ from what previously has been seen. In this work, we develop and evaluate methods for a Mars rover to use *novelty* to guide the selection of observation targets with the goal of accelerating discovery. In a study comparing three image content representations and five novelty-based ranking methods, we found that the Isolation Forest identified the largest number of novel targets using a combination of intensity and shape features to represent the candidate targets. It was followed closely by the Local RX algorithm using raw pixel features. All algorithms achieved performance well above alternatives such as random selection or selecting the best match to current science objectives, which do not account for novelty.

1 INTRODUCTION

The Mars Science Laboratory rover (Curiosity) has been exploring the surface of Mars since August 2012. The primary goal of the mission is to determine whether Mars may have been habitable in the past (or present) [6]. The Mars 2020 rover (Perseverance) launched in July 2020 and will arrive at Mars in February 2021. The Perseverance mission will additionally search for signs of past microbial life itself and cache promising samples for a future mission to recover and return to the Earth for further study [13].

Mars rover surface operations are planned by a team of scientists and engineers on the Earth. After reviewing the rover's latest position, status, and the data that it collected, the team develops a plan for the rover's next day that specifies the time and duration of each activity: drives, data collection, communication passes, etc. The plan is carefully simulated and tested to confirm that it fits within the available resources (time, energy, viewing opportunities) before it is uplinked to the rover.

At the end of each drive, the Mars Science Laboratory (MSL) rover collects images of its new location to transmit to Earth as an aid to planning the next day's investigation. In 2016, the rover gained the ability to autonomously select where to point the ChemCam instrument to acquire compositional spectra for targets that match specified science priorities. ChemCam employs a laser-induced breakdown spectrometer to fire a laser and then record the resulting plasma spectrum which reveals the presence of individual elements [12]. ChemCam can reach targets up to seven meters from the rover without requiring a physical approach and instrument contact.

The software onboard the rover that enables decisions for autonomous targeting is AEGIS (Autonomous Exploration for Gathering Increased Science) [5]. It was originally developed for the Mars Exploration Rovers to detect interesting targets in navigation cameras for follow-up observations with the higher resolution, multi-band panoramic camera [3]. For MSL, target detection and ranking is likewise done using navigation camera images, and the top-ranking target's location is used to command the collection of spectra by ChemCam. The science team specifies a target "signature" that expresses the current science priorities in terms of target size, shape, smoothness, albedo (average pixel intensity), etc. AEGIS has autonomously selected over 250 targets for MSL since 2016 and has led to increased use of ChemCam by the science team overall [5].

In this paper, we propose and evaluate a new approach to autonomous targeting that can allow Mars rovers to select targets in terms of novelty. This capability can complement the AEGIS system by identifying additional targets that may not match current science objectives but which present the possibility of a new discovery. We compared multiple algorithms for novelty-based ranking on a collection of Mars rover images and assesses how well their selections agreed with an independent manual identification of novel targets. We found that the Isolation Forest was most successful at discovering novel targets when using features that cap-

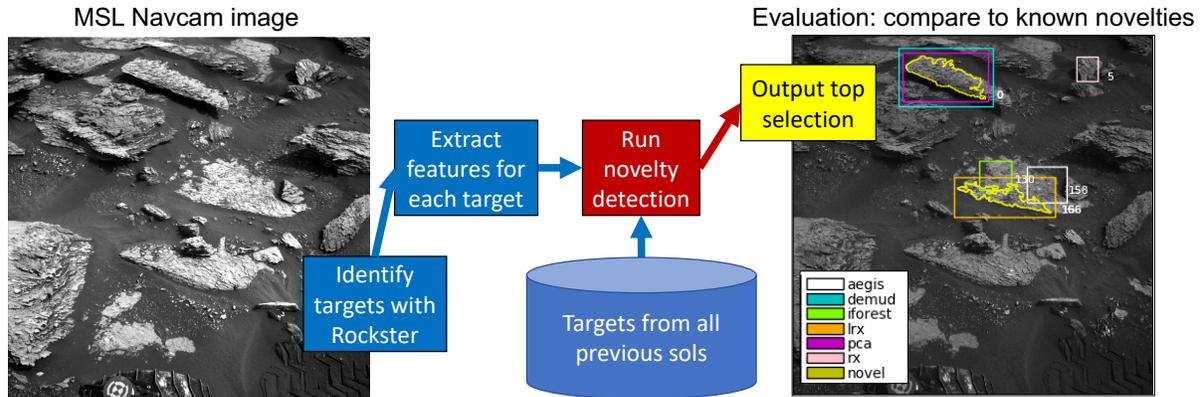


Figure 1: Novelty-based target ranking identifies targets, extracts features to represent each target, uses novelty detection to rank the targets in comparison to previously observed targets, and outputs the top choice. Evaluation is conducted by comparing selections to manually identified novel targets (gold polygons). Example image was collected by MSL on sol 1683.

ture both intensity and shape attributes. The Local RX algorithm achieved almost the same level of performance using the raw pixel features. All five algorithms performed well above random selection as well as the use of AEGIS alone, which was not designed to prioritize novel targets.

2 NOVELTY-BASED TARGETING FOR MARS ROVERS

We propose a new capability for the onboard, autonomous selection of follow-up targets by Mars rovers. In addition to the ability to select targets that match pre-defined signatures, we add the option to select targets based on their novelty with respect to what the rover has already encountered.

Fig. 1 illustrates the steps in this procedure applied to a navigational camera (Navcam) image taken by the MSL rover on sol (mission day) 1683. Candidate targets are detected using the Rockster [2] image analysis system, which is part of AEGIS. Rockster segments the image into individual targets using edge detection and grouping. It was designed to run efficiently enough, in terms of runtime and memory consumption, for operation onboard Mars rovers. The Mars Science Laboratory rover has a processor that runs at 133 MHz, with only 16 megabytes of memory available for AEGIS. The result of Rockster is a collection of polygons that describe the outline of each target.

Each target is represented using a standard set of features, as described in section 2.1. Next, a novelty detection method is used to rank the targets in the current image in terms of their novelty with respect to the collection of targets previously observed on the mission (section 2.2). Finally, the system outputs the top selection, which can be used to inform further follow-up

data collection by ChemCam or other instruments on the rover.

2.1 Target Representation

We investigated three methods for representing the target image content to provide input for the novelty ranking algorithms. Each representation converts data from the two-dimensional grayscale image that contains the target into a one-dimensional feature vector. Targets vary in size and aspect ratio. The most basic “intensity_pixels” representation uses the pixel intensity values directly as features. To provide a consistent dimensionality despite different target sizes, we first crop a square bounding box around each target polygon and resize the cropped image to 64×64 pixels. The cropped image containing the target is converted into a feature vector with 4096 values.

Second, the “intensity_stats” representation computes seven statistical features to summarize the contents of the pixels within the target image. These statistics include the minimum, maximum, mean, and median pixel values as well as their standard deviation, skew, and kurtosis. Since these aggregate statistics abstract away from the original pixels, we do not need to resize the target images to a standard size, so they operate directly on the cropped (bounding box) image. The result is a 7-dimensional feature vector for each target.

Finally, the “intensity+shape” representation uses nine domain-specific features computed by the AEGIS system for its target ranking procedure. This representation makes use of the polygon itself as well as the pixel values. Pixel features include the mean and standard deviation of pixels within the polygon (not the entire bounding box). The geometry of the target is captured by the area, perimeter, and “ruggedness” of the polygon

as well as features derived from fitting an ellipse to the polygon and recording its eccentricity, orientation, and semi-major and semi-minor axis lengths [4].

2.2 Novelty Ranking Algorithms

We implemented several methods for assigning a novelty score to each target, including ranking the targets by reconstruction error when using a lower-dimensional model computed with Principal Component Analysis (PCA) or DEMUD [11], by sparsity in feature space (Isolation Forest [8]), and by deviation from either the global “background” (RX) [10, 1] or a local window (LRX).

One approach to novelty detection is to compare new observations to a compact model of content from prior observations. Any new content that is not recognized by the model provides a potential indication of novelty. We quantify this kind of novelty as the *reconstruction error* between the new observation and its reconstruction using the model. For example, we use Principal Component Analysis (PCA) to compute k eigenvectors of a $n \times d$ training data set, project each item into the corresponding k -dimensional space, and then reconstruct the original observation in d -dimensional space. The PCA-based novelty score for a feature vector \mathbf{x} is:

$$s_{PCA}(\mathbf{x}) = \|\mathbf{x} - (\mathbf{U}\mathbf{U}^T(\mathbf{x} - \mu) + \mu)\|_2 \quad (1)$$

where \mathbf{U} contains the eigenvectors and μ is the mean feature vector. Since the projection onto \mathbf{U} discards information, some details may be lost in the reconstruction, and these increase the reconstruction error score. The DEMUD algorithm [11] uses the same scoring method, but it updates \mathbf{U} to include each new selection once it is made, to reduce the chance of a redundant selection. We used a k value of 10 for “intensity_pixels”, 7 for “intensity_stats”, and 3 for “intensity+shape” features.

In contrast, the Isolation Forest [8] identifies items that are easily separated from the majority of the data set. An Isolation Forest consists of multiple binary trees, in which each tree employs a series of randomly chosen splits to partition the feature space. Each split specifies a randomly chosen feature and random threshold value. The position of each item in feature space dictates its outlier (novelty) score, which is proportional to the number of splits required to “isolate” the item into its own cell in feature space. Outliers will tend to have a lower score, while inliners will have a higher score. The final isolation forest score for an item \mathbf{x} is its average score across all trees in the forest. We used the default value of 100 trees for the Isolation Forest.

The Reed-Xiaoli (RX) detector [10] is an unsupervised anomaly detection method that has been used successfully in remote sensing and exploration applications (e.g., [1, 7, 14]). RX computes an anomaly score for each item \mathbf{x} using the Mahalanobis distance between the item and a background distribution:

$$s_{RX}(\mathbf{x}) = (\mathbf{x} - \mu)^T \Sigma^{-1} (\mathbf{x} - \mu) \quad (2)$$

where Σ is the covariance matrix of the “background”. We used a separate training set as the background for RX.

Since we are working with image data, we can also use a variant of RX called Local RX (LRX), which generates a score for each pixel, then averages those scores to generate a score for the full image. LRX computes μ and Σ using a pair of sliding windows (inner and outer) to define the local “background” around each pixel as the pixels that fall between the inner and outer windows. LRX has been widely used to detect small anomalies in hyperspectral images [9]. In our experiments, we used an inner window size of 3×3 pixels and the outer window size was 5×5 pixels.

3 EXPERIMENTAL RESULTS

3.1 Data Set

When AEGIS is used for MSL, it operates on a pre-planned Navigation camera (Navcam) grayscale image that is collected by the rover. The Rockster [2] algorithm identifies candidate targets within the Navcam image. The Navcam image ID, Rockster-identified targets, and target prioritization signatures (which express the current science priorities and are used to rank candidate targets) are included as part of the information sent from the rover to the planning team. AEGIS was first deployed on sol 1343 (May 16, 2016), and our data consists of the Navcam images and AEGIS targets identified from sol 1343 to 2578 (November 6, 2019). We assembled a data set of 6,005 targets that had an area of at least 100 pixels. We accessed the image data directly from the Navcam Experiment Data Records (EDRs) from the Planetary Data System (PDS) Imaging Node (https://pds-imaging.jpl.nasa.gov/data/msl/MSLNAV_0XXX/). The EDRs are stored as 16-bit images, which we converted to 8-bit images by multiplying all pixels by a factor of 255/4095, which is the same conversion done onboard MSL for AEGIS.

To create a reference dataset of “novel” targets for evaluating algorithm ranking performance, co-author R. Francis identified which of the Rockster targets might have been considered novel or interesting by the MSL science team and had not been observed in Navcam im-

ages on prior sols. For example, Fig. 1 shows two targets identified as novel on sol 1683 because they depict plate-like exposures of eroded layers. Other kinds of novel target features include unusual dark-toned material, light-toned veins, and rough surface textures. In all we obtained a list of 108 reference novel targets observed from sols 1343 to 1703.

For the image shown in Fig. 1, Rockster identified a total of 11 targets, two of which are novel (gold polygons). The DEMUD and PCA algorithms selected target 0 (novel), and the LRX algorithm selected 166 (novel). The Isolation Forest selected target 130, and RX selected target 5, neither of which were considered novel. The target selected by AEGIS (158) was also not considered novel, which is unsurprising since AEGIS has a different objective.

3.2 Methodology

We performed a series of experiments to evaluate the ability of each algorithm to identify novel targets. Within our novel target reference data set, we identified 28 sols that had at least one novel reference target. For each of these sols, we ranked all candidate targets identified by Rockster in the Navcam image using the novelty ranking algorithms described in Section 2.2. We defined a “prior” (training) data set of targets for each algorithm that consisted of all targets observed (novel or not) on sols prior to the experiment sol, e.g., sols 1343 to $S-1$ where S is the experiment sol. This simulates the operational scenario onboard the rover in which the novelty ranking algorithm would rank new targets based only on data from previous sols as it traverses the Mars surface. We evaluated each algorithm with the three target representations (“intensity_pixels”, “intensity_stats”, and “intensity+shape”) described in Section 2.1. Each ranking algorithm provides a novelty score for each target, with targets ranked in order from most novel to least novel.

Some algorithms restrict which feature representations they can employ. Since LRX operates on local windows surrounding each pixel in the target image, it cannot be used with the “intensity_stats” or “intensity+shape” features. RX requires that the number of items n be greater than the number of features d , because it computes a feature covariance matrix (Equation 2) that is singular if $n < d$. Since each target has dimension 64×64 , the number of features for the “intensity_pixels” representation is 4096. For all sols in our experiments, the number of prior images is less than 4096, so we did not evaluate RX with the “intensity_pixels” representation.

When AEGIS is used onboard the MSL rover, typi-

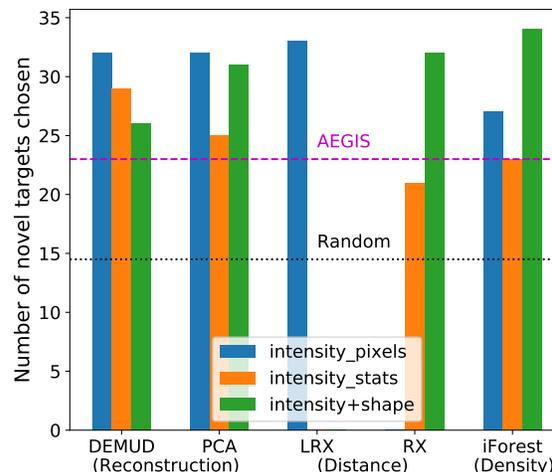


Figure 2: Summary of novelty detection performance using each representation. Dotted lines show Random and AEGIS algorithm performance.

cally only one or two targets are chosen by AEGIS for follow-up observation with ChemCam due to rover resource constraints. To quantify the performance of the novelty ranking algorithms in a similar operational scenario, we recorded how many novel reference targets were found in the top two ranked targets for each algorithm (“top-2 score”). Although AEGIS was not designed as a novelty detection algorithm, we also computed its performance on the same task for comparison. As a baseline, we implemented a random ranking algorithm that employs a pseudo-random number generator to randomly rank the targets in each experiment. We ran the random algorithm with 10 different random seeds and reported the average performance.

3.3 Results

The top-2 score for each novelty ranking algorithm, summed across all scenarios, is shown in Fig. 2. Individual results for each sol are given in the Appendix. The AEGIS and random algorithms do not employ different feature representations, so they each have a single total top-2 score of 23 (AEGIS) and 14.5 (random, 10 trials), respectively. The maximum possible score is 54.

The best top-2 score (34) was achieved by the Isolation Forest using the “intensity+shape” features (Fig. 2). LRX, DEMUD, and PCA achieved their best scores with the raw pixel intensities (33, 32, and 32 respectively). The best RX result (32) was achieved with the “intensity+shape” features. Overall, the best score achieved by each algorithm was well above that of random selection. All novelty algorithms, using their individual best feature representations, also out-performed

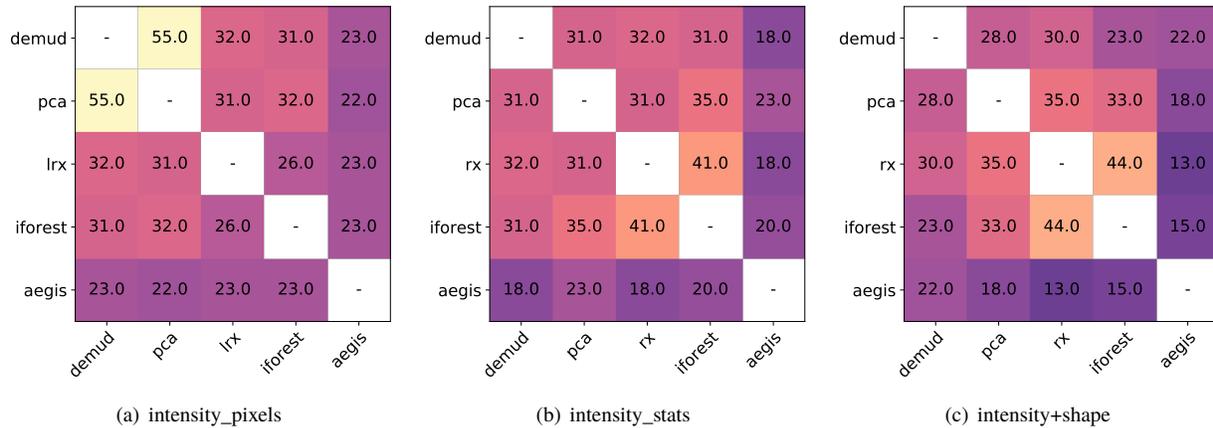


Figure 3: Agreement between algorithms in terms of the number of common selections within their top two choices

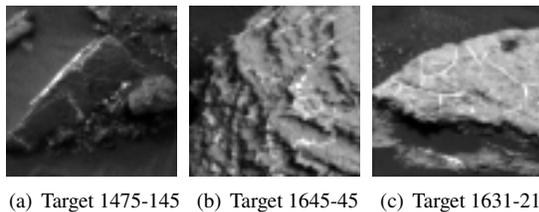


Figure 4: Targets found uniquely by one algorithm (see text)

the AEGIS algorithm. AEGIS was not designed for selecting novel targets; its utility in this comparison is to provide insight into how the current onboard system would perform on the task of novelty detection. As we hypothesized, using an algorithm designed to rank items by their novelty achieves the best result.

4 DISCUSSION

4.1 Feature representations

Compared to algorithms that assign a score to each target based on its novelty with respect to targets in the prior data set, LRX scores each pixel in the target image based on its novelty with respect to the local pixels surrounding it. This makes LRX effective for identifying small features within candidate target images that deviate from their surrounding patterns, such as light-toned veins. These veins are rare within the full target set, but they are often selected by humans for inclusion in the novel reference target data set. For example, novel target 145 from sol 1475 in Fig. 4(a), which contains a small bright vein, was found only by the LRX method and no other algorithm when using “intensity_pixels” features.

The Isolation Forest had the lowest performance using raw pixel intensities but the highest score with the

“intensity+shape” features. The Isolation Forest “isolates” candidate targets by recursively selecting one of the features at random and then partitioning targets by randomly selecting a split value between the maximum and minimum values of the selected feature. This process is represented as a tree structure, and the number of splits (branches) required to isolate a candidate target (averaged over many random trees) is the Isolation Forest’s measure of novelty—i.e., fewer splits implies a higher degree of novelty. Isolating targets in this manner using a large number of features as in the case of raw pixel intensities is likely less effective than when fewer, more descriptive features such as the “intensity+shape” features are used. Examples of two targets found only by the Isolation Forest method when using the latter features are shown in Fig. 4(b) and Fig. 4(c); both targets contain multiple veins.

4.2 Algorithm agreement

In addition to reporting the agreement between algorithm selections and novel reference target selections, it is also useful to assess agreement between algorithms themselves. For example, two algorithms that have high performance but make different selections could be used in complement onboard to promote a diversity of novel target selections. Figs. 3(a)-3(c) show agreement between algorithms for the experiments in Section 3.3. In each matrix, the entries represent the number of common targets in the selections made by each pair of algorithms out of 56 total selections. We omitted the diagonal entries because these represent the agreement of each algorithm with itself. DEMUD and PCA using “intensity_pixels” features (Fig. 3(a)), stand out for agreeing on nearly all selections (55 out of 56). Both algorithms rank the images based on reconstruction error using $k = 10$ eigenvectors to model variation within

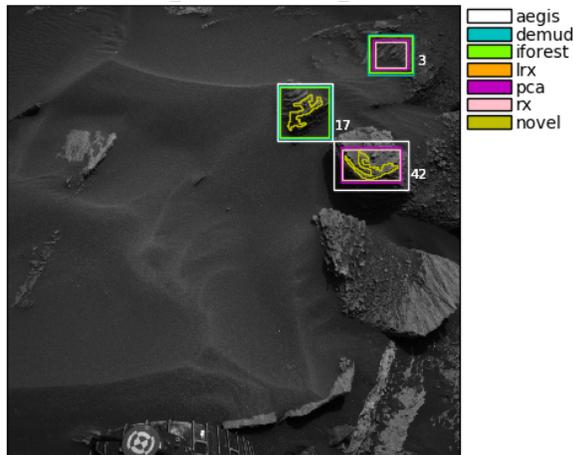


Figure 5: Top 2 selections by novelty ranking algorithms on sol 1685 (“intensity+shape” features)

the data set. The reduction from 4096 to 10 features yields similar results with the two methods.

For “intensity_stats” (Fig. 3(b)) and “intensity+shape” (Fig. 3(c)) features, we see generally lower agreement (more diversity). RX and iForest are the algorithms with the highest agreement in these settings, so it seems likely that distance from “prior” targets (RX) is correlated with separability (Isolation Forest) in these representations. Using “intensity_pixel” features, PCA, DEMUD, and LRX achieved similarly high top-2 scores (Fig. 2), but LRX had lower agreement with PCA (31) and DEMUD (32), indicating that its selections complement theirs. LRX selected 9 targets not chosen by PCA or DEMUD, of which three contained veins, two had rough surface texture, and four showed layering. PCA and DEMUD selected 6 targets not chosen by LRX, of which only one contained a vein. Likewise, using “intensity+shape” features, PCA, iForest, and RX had similarly high top-2 scores, but PCA made the same selections as iForest and RX in only 31 to 35 cases.

4.3 Novel reference targets

We observed that some targets were selected by multiple novelty ranking algorithms that were not in the novel reference target list. We investigated these choices carefully and found that they revealed criteria that science team members factor into target prioritization in addition to scientific interest or novelty. For example, Fig. 5 shows the top 2 selections made by AEGIS and the novelty ranking algorithms using the “intensity+shape” features. The novel reference targets for this experiment were 17 and 42, chosen because they contain plate-like exposures of eroded layers. While most of the ranking algorithms selected

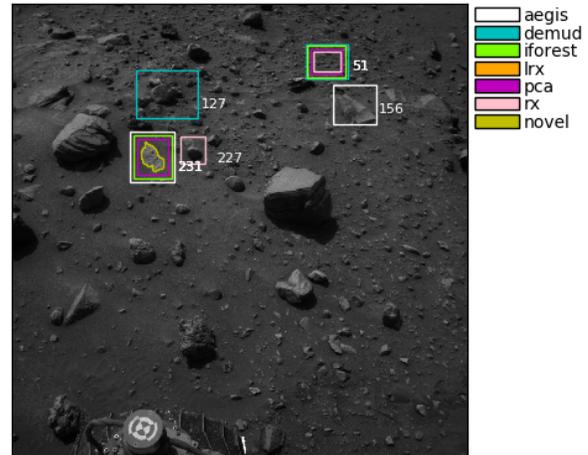


Figure 6: Top 2 selections by novelty ranking algorithms on sol 1400 (“intensity_stats” features)

these targets, some also selected target 3. While target 3 contains the same plate-like exposures as targets 17 and 42, target 3 is farther away from the rover and is partially covered by sand. While target 3 is a valid choice from a novelty perspective, this target is not as desirable as 17 and 42 because the follow-up Chem-Cam measurement of target 3 would likely hit the sand rather than the rock itself, and a measurement of the rock composition is generally more desirable to the science team.

Fig. 6 shows another experiment from sol 1400 using “intensity_stats” features. In this experiment, target 231 was the novel reference target, desirable because it is bedrock that is notably light-toned. DEMUD, PCA, and AEGIS did select target 231, but the algorithms also prioritized other targets (51, 156, 227, and 127). These other targets, originally identified as candidate targets by the Rockster algorithm, appear to merge multiple rocks into a single target and as a result might have uncommon feature values that cause them to be prioritized by the novelty ranking algorithms.

5 CONCLUSIONS AND FUTURE WORK

We propose a new capability for onboard data analysis for Mars rovers in which they can assess the novelty of each candidate target and use this information to inform autonomous decisions about which targets merit additional follow-up study. Novelty-based targeting would not replace the method currently used by AEGIS onboard MSL to identify targets that best match current science objectives. Instead, mission scientists and planners could selectively employ novelty-based targeting to complement the AEGIS selections (e.g., one science-based and one novelty-based target per drive). In addition, in cases where resources allow for only a single

Table 1: Number of novel reference targets found in each algorithm’s top two selections, with “intensity_pixels” features

sol	DEMUD	iForest	LRX	PCA
1347	1	1	1	1
1371	1	1	0	1
1400	0	0	0	0
1439	2	1	1	2
1469	0	0	0	0
1475	1	0	2	1
1519	0	0	0	0
1521	2	1	2	2
1583	0	0	0	0
1605	2	1	2	2
1612	0	1	1	0
1629	2	2	2	2
1631	1	1	1	1
1636	2	2	1	2
1645	2	2	2	2
1660	2	2	2	2
1666	1	1	1	1
1672	2	2	2	2
1673	2	0	0	2
1676	0	0	1	0
1683	1	0	2	1
1684	1	0	1	1
1685	2	2	2	2
1686	2	1	2	2
1690	1	1	2	1
1697	0	1	1	0
1699	1	2	1	1
1703	1	2	1	1
Total (of 54)	32	27	33	32

follow-up target to be chosen, planners could specify a threshold on the novelty score, rather than a simple ranking, to determine when the rover has encountered a target of sufficient novelty to supersede the target that best matches current science priorities.

An important consideration for any use of novelty-based ranking onboard a Mars rover is the resources (memory and computation) that it requires. Our next step will be to assess each algorithm in terms of runtime and memory consumption to determine how it would operate given only 16 MB of RAM and a RAD750 processor (133 MHz). In addition, we will evaluate these methods to rank targets identified in color images such as those collected by the Mastcam instrument on MSL. Previous work has shown that spectral information is leveraged in different ways by these algorithms [7], and the best algorithm choice may be different when color information is available.

Appendix A: Per-sol Experiment Results

Performance for each novelty ranking algorithm for each sol scenario we evaluated is given in Tab. 1 (“intensity_pixels”), Tab. 2 (“intensity_stats”), and Tab. 3 (“intensity+shape”). The bottom row of each table

Table 2: Number of novel reference targets found in each algorithm’s top two selections, with “intensity_stats” features

sol	DEMUD	iForest	PCA	RX
1347	1	1	1	1
1371	1	1	1	0
1400	0	1	1	0
1439	1	2	1	1
1469	1	0	0	0
1475	1	1	2	1
1519	0	0	0	0
1521	2	1	2	1
1583	1	1	1	1
1605	2	2	2	2
1612	1	1	0	0
1629	2	2	2	2
1631	2	2	2	2
1636	2	1	1	1
1645	0	0	0	0
1660	2	1	2	1
1666	0	0	0	0
1672	2	2	2	2
1673	0	0	0	0
1676	0	0	0	0
1683	1	0	0	1
1684	1	0	1	0
1685	1	1	1	1
1686	1	1	2	2
1690	0	0	0	0
1697	1	1	0	0
1699	2	1	1	1
1703	1	0	0	1
Total (of 54)	29	23	25	21

gives the total score across all sols. The AEGIS and random algorithms do not employ different feature representations, so they each have a single total top-2 score of 23 (AEGIS) and 14.5 (random, 10 trials), respectively. The maximum possible score is 54.

Acknowledgments

We thank the NASA Center Innovation Fund for supporting this work. Part of this research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. Copyright 2020. All rights reserved.

References

- [1] Burl, M. C. and Lucchetti, D. (2000). Autonomous visual discovery. In *Data Mining and Knowledge Discovery: Theory, Tools, and Technology II*, pages 240–248.
- [2] Burl, M. C., Thompson, D. R., deGranville, C., and Bornstein, B. J. (2016). Rockster: Onboard rock segmentation through edge regrouping. *Journal of Aerospace Information Systems*, 13:329–342.
- [3] Estlin, T., Bornstein, B., Gaines, D., Anderson, R., Thompson, D., Burl, M., Castaño, R., and Judd, M. (2012). AEGIS automated science targeting for the MER

Table 3: Number of novel reference targets found in each algorithm's top two selections, with "intensity+shape" features

sol	DEMUD	iForest	PCA	RX
1347	1	1	1	0
1371	0	0	0	0
1400	0	0	0	0
1439	1	1	2	0
1469	1	1	1	1
1475	0	2	1	1
1519	0	0	0	0
1521	2	2	1	2
1583	1	1	1	1
1605	1	1	1	1
1612	1	1	1	1
1629	1	2	2	2
1631	1	1	1	2
1636	1	1	2	2
1645	2	2	2	2
1660	1	1	1	1
1666	2	2	2	2
1672	2	2	2	2
1673	0	2	0	1
1676	0	1	1	1
1683	1	2	2	2
1684	1	1	0	1
1685	1	2	1	2
1686	2	2	2	2
1690	1	1	1	1
1697	0	0	0	0
1699	2	1	2	1
1703	0	1	1	1
Total (of 54)	26	34	31	32

Opportunity rover. *ACM Transactions on Intelligent Systems and Technology*, 3(3):Article 50.

- [4] Estlin, T., Gaines, D., Bornstein, B., Schaffer, S., Tompkins, V., Thompson, D. R., Altinok, A., Anderson, R. C., Burl, M., no, R. C., Blaney, D., Flores, L. D., Nelson, T., and Wiens, R. (2014). Automated targeting for the MSL rover ChemCam spectrometer. In *Proceedings of the 12th International Symposium on Artificial Intelligence, Robotics, and Automation in Space (i-SAIRAS)*.
- [5] Francis, R., Estlin, T., Doran, G., Johnstone, S., Gaines, D., Verma, V., Burl, M., Frydenvang, J., Montañó, S., Wiens, R. C., Schaffer, S., Gasnault, O., DeFlores, L., Blaney, D., and Bornstein, B. (2017). AEGIS autonomous targeting for ChemCam on Mars Science Laboratory: Deployment and results of initial science team use. *Science Robotics*, 2(7).
- [6] Grotzinger, J. P., Crisp, J., Vasavada, A. R., Anderson, R. C., Baker, C. J., Barry, R., Blake, D. F., Conrad, P., Edgett, K. S., Ferdowski, B., Gellert, R., Gilbert, J. B., Golombek, M., Gómez-Elvira, J., Hassler, D. M., Jandura, L., Litvak, M., Mahaffy, P., Maki, J., Meyer, M., Malin, M. C., Mitrofanov, I., Simmonds, J. J., Vaniman, D., Welch, R. V., and Wiens, R. C. (2012). Mars Science Laboratory mission and science investigation. *Space Science Review*, 170:5–56.
- [7] Kerner, H. R., Wagstaff, K. L., Bue, B. D., Wellington, D. F., Jacob, S., Horton, P., III, J. F. B., Kwan, C., and Amor, H. B. (2020). Comparison of novelty detection methods for multispectral images in rover-based planetary exploration missions. *Data Mining and Knowledge Discovery*.
- [8] Liu, F. T., Ting, K. M., and Zhou, Z.-H. (2008). Isolation forest. In *Proceedings of the IEEE International Conference on Data Mining*, pages 413–422.
- [9] Molero, J. M., Garzón, E. M., García, I., and Plaza, A. (2013). Analysis and optimizations of global and local versions of the RX algorithm for anomaly detection in hyperspectral data. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 6(2):801–814.
- [10] Reed, I. S. and Yu, X. (1990). Adaptive multiple-band CFAR detection of an optical pattern with unknown spectral distribution. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 38(10):1760–1770.
- [11] Wagstaff, K. L., Lanza, N. L., Thompson, D. R., Dieterich, T. G., and Gilmore, M. S. (2013). Guiding scientific discovery with explanations using DEMUD. In *Proceedings of the Twenty-Seventh Conference on Artificial Intelligence*, pages 905–911.
- [12] Wiens, R. C., Maurice, S., and the ChemCam team (2011). The ChemCam instrument suite on the Mars Science Laboratory rover Curiosity: Remote sensing by laser-induced plasmas. *Geochemical News*, 145.
- [13] Williford, K. H., Farley, K. A., Stack, K. M., Allwood, A. C., Beaty, D., Beegle, L. W., Bhartia, R., de la Torre Juarez, A. J. B. M., Hamran, S.-E., Hecht, M. H., Hurowitz, J. A., Rodriguez-Manfredi, J. A., Maurice, S., Milkovich, S., and Wiens, R. C. (2018). *From Habitability to Life on Mars*, chapter 11: The NASA Mars 2020 Rover Mission and the Search for Extraterrestrial Life, pages 275–308. Elsevier.
- [14] Zhou, J., Kwan, C., Ayhan, B., and Eismann, M. T. (2016). A novel cluster kernel RX algorithm for anomaly and change detection using hyperspectral images. *IEEE Transactions on Geoscience and Remote Sensing*, 54(11):6497–6504.