COMPARISON OF NOVELTY DETECTION METHODS FOR MULTISPECTRAL IMAGES FROM THE MASTCAM INSTRUMENT ONBOARD MARS SCIENCE LABORATORY.

H. R. Kerner[1], K. L. Wagstaff[2], Brian Bue[2], Danika F. Wellington[1], Samantha Jacob[1], James F. Bell III[1], and Heni Ben Amor[3], [1]School of Earth and Space Exploration, Arizona State University, Tempe, AZ, hkerner@asu.edu, [2]Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, [3]School of Computing, Information, and Decision Systems Engineering, Arizona State University, Tempe, AZ.

Introduction: Rover missions like the Mars Science Laboratory (MSL) operate through a system of tactical planning: the rover is commanded to make observations on the surface of Mars and sends the science data back to Earth, the science team analyzes the latest data and decides what observations to make on the next sol based on that data, and the process repeats. There is limited time available for science planning: <12 hours for MSL and potentially as few as 5 hours for the upcoming Mars 2020 rover. Thus, there is a need for systems that can rapidly and intelligently extract information of interest from science instrument data to focus on potential discoveries and avoid missed science opportunities. Science planning and data analysis teams could benefit by spending their limited available time on the most promising, or novel, observations. The goal of novelty detection (also referred to as anomaly or outlier detection) techniques is to identify patterns in data that have not been previously or frequently observed [1,2]. We evaluated multiple methods for detecting novel geology in multispectral images acquired by MSL.

Mastcam Multispectral Images: One instrument the MSL rover uses to make geologic observations is the mast camera, or "Mastcam," a pair of CCD imagers mounted on the rover's mast ~2 meters above the surface [3-5]. Each of Mastcam's cameras, or "eyes," has an eight-position filter wheel enabling images to be acquired in "true color" (Bayer pattern broadband red, green, and blue) and with six narrowband spectral filters spanning ~400-1100 nm (visible to short-wave near-infrared) [3]. A similar camera, Mastcam-Z, will be onboard the Mars 2020 rover [6]. Examples of novel geology in Mastcam images include iron meteorites [7-8] and broken rocks that expose mineralogy under the dusty surface (*e.g.*, **Fig. 1**).

We created a dataset using 477 multispectral thumbnails (smaller, immediately-downlinked versions of full-resolution images) acquired between MSL sols 1-1666 using the right eye (M-100). From these, we identified 156 images with novel geology and cropped 237 64x64x6-pixel sub-frames around the novel regions (*e.g.,* **Fig. 1**). We considered all other images *typical* and divided them into training, validation, and test sets. We augmented these datasets by sub-sampling 64x64x6-pixel sub-frames with a sliding window (9,302 train, 1,386 validation, and 1,302 test sub-frames). We used the combined 237 novel and 1,302 typical test images for testing the novelty detection methods.

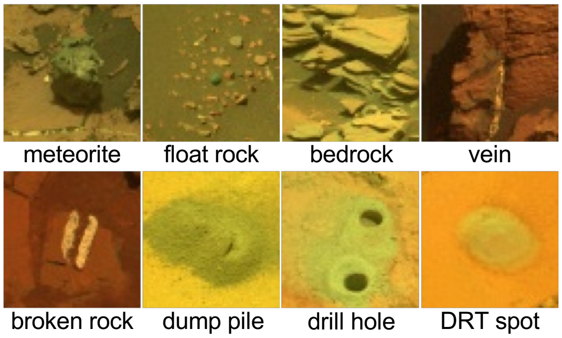
meteorite float rock bedrock vein
broken rock dump pile drill hole DRT spot

Figure 1 Examples of novel geology in Mastcam images.

Methods: We evaluated two traditional methods for novelty detection: Reed-Xiaoli (RX) detectors [9] and singular value decomposition (SVD) [10]. We chose these methods because they are well established and enable visualization of detections within images. We also evaluated two recent deep learning methods: convolutional autoencoders (CAE) [11] and generative adversarial networks (GANs) [12]. Compared to traditional methods, these methods often exhibit better performance for high-dimensional image datasets.

RX Detector: The RX method is commonly used to detect anomalous pixels in multispectral or hyperspectral images. RX assigns an anomaly score to each pixel that is the Mahalanobis distance between the pixel and a background distribution [9]. The background distribution is usually defined as all other pixels or a window of pixels in the image. We instead computed the RX score for each pixel with respect to the background of the entire *typical* training dataset.

SVD: SVD is a method of computing the principal components of a dataset. The principal components are the eigenvectors of the data matrix [1], and the reconstruction error of an input can be interpreted as a novelty score [10]. We used the first 100 principal components for modeling the *typical* training dataset.

CAE: An autoencoder is a type of neural network that learns the salient features in a dataset. A convolutional autoencoder (CAE) consists of an *encoder* network to map (compress) inputs to a low-dimensional encoding and a *decoder* network to reconstruct inputs from the encoding using convolutional layers [11]. In training, the loss between input and reconstruction is minimized. Loss is usually defined to be the mean squared error (MSE). We used this approach in [13] and found that MSE loss resulted in noisy reconstructions that lead to higher residuals for non-novel regions and thus false positives. To improve the quality of autoencoder reconstructions, we propose to maximize the structural similarity index (SSIM) while minimizing MSE during training. SSIM is a measure of *perceived* similarity proposed by Wang *et al.* (2010) [14] for evaluating image compression algorithms (*e.g.*, JPEG). We trained a CAE with the same architecture as in [13] using this hybrid loss function and Euclidean distance between inputs and reconstructions as the novelty score for test examples.

GAN: GANs are a type of neural network that learn data-generating distributions for a dataset via minimax optimization of two networks. The *generator* samples from a *d*-dimensional normal distribution (*d* is size of latent vector) and tries to reconstruct an image resembling training images. The *discriminator* tries to distinguish training images from generated images. The discriminator minimizes its loss while the generator maximizes discriminator loss. As in [12], we used a bi-directional GAN (BiGAN [15]) to simultaneously train a third *encoder* network, which maps training images to *d*-dimensional encodings. We trained a BiGAN on the *typical* dataset and computed novelty scores as a weighted sum of the discriminator loss and reconstruction error as in [12].

Experimental Results: We evaluated each novelty detection method on the combined *novel* and *typical* test dataset. We computed the AUC (area under the curve) of the Receiver Operating Characteristic (ROC) curve and precision at N (P@N) [16] to compare performance. P@N is the fraction of images from the *novel* dataset included in the top N selections by novelty score. To assess model performance on different types of novel geology in Mastcam images, we divided the *novel* dataset into 8 sub-classes: meteorite, bedrock, broken rock, float rock, vein, dust removal tool (DRT) spot, drill hole, and dump pile. **Fig. 2** shows ROC AUC scores for the novel dataset overall and individual sub-classes. **Table 1** gives P@N.

We found the CAE had the best P@N performance and the methods had comparable AUC score performance overall. However, we found significant differences for different types of novel geology. We found that SVD had the best performance in the bedrock, broken rock, float, and meteorite classes and was closely followed by the CAE. Because SVD and CAE minimize reconstruction error during training, they may learn mappings that reconstruct most inputs reasonably well even if the image contains unseen features like drill holes, DRT spots, and dump piles. The GAN had the best performance in the drill hole, DRT, and dump pile classes. This may be because the GAN discriminator learns to detect features common in *typical* images during training, which might not be recognized in the drill hole, DRT spot, and dump pile classes. RX had the worst performance overall by both measures; inspecting false positives suggests RX is sensitive to color discontinuities caused by JPEG artifacts, which is problematic for thumbnail images.

Our experiments suggest multiple methods may be needed to detect the diverse novel geology observed by Mastcam along MSL's traverse. In future work, we will perform additional experiments to constrain the performance of these and other novelty detection methods for Mastcam multispectral images. Additionally, we are exploring ways to incorporate these methods into operations for MSL as well as the search capabilities of the Planetary Data System (PDS).

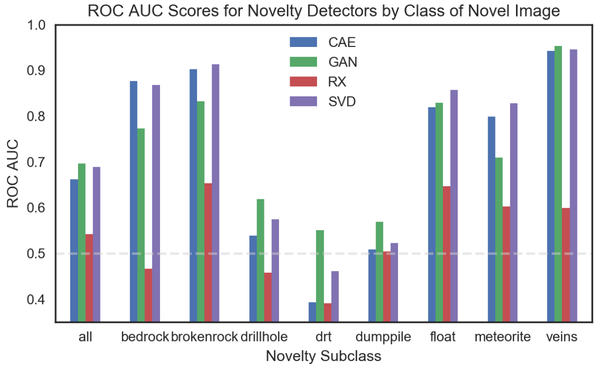

Figure 2 ROC AUC scores for *typical* test dataset and *novel* dataset.

Table 1 Precision at N (fraction of *novel* images in top N selections).

Method	N=15	N=50	N=200
RX	0.27	0.14	0.11
SVD	0.60	0.38	**0.36**
CAE	**0.87**	**0.52**	0.32
GAN	0.20	0.38	0.33

Acknowledgments: This research was carried out (in part) at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration and funded through the Internal Strategic University Research Partnerships (SURP) program.

References: [1] Markou M. et al. (2003) *Signal Processing*, 83(12), 2481-2497. [2] Markou M. et al. (2003) *Signal Processing*, 83(12), 2499-2521. [3] Bell III J. F. et al. (2017) *ESS,* 4(7), 396-452. [4] Malin M. et al. (2017) *ESS*, 4(8), 506-539. [5] Grotzinger J. et al. (2012) *SSR*, 170(1-4), 5-56. [6] Bell III J. F. (2016) *Intl Workshop on Instrumentation for Planetary Missions.* [7] Wellington D. et al. (2017) *American Mineralogist*, 102(6), 1202-1217. [8] Johnson J. et al. (2014) *AGU Fall Meeting*. [9] Reed I. S. & Yu X. (1990) *IEEE Trans Signal Proc*, 38(10), 1760-1770. [10] Wagstaff K. L. et al. (2013) *AAAI*, 905-911. [11] Masci J. et al. (2011) *ICANN,* 52-59. [12] Zenati H. et al. (2018) *ICDM,* 727-736. [13] Kerner H. R. et al. (2019) *IAAI.* [14] Wang Z. et al. (2004) *IEEE Trans Image Proc*, 13(4), 600-612. [15] Donahue J. et al. (2017) *ICLR*. [16] Campos G. et al. (2016) *Data Mining & Knowledge Discovery*, 30(4), 891-927.