

Novelty Detection for Multispectral Images with Application to Planetary Exploration

Hannah R Kerner¹, Danika F Wellington¹, Kiri L Wagstaff³, James F Bell¹, Heni Ben Amor²

¹School of Earth and Space Exploration, Arizona State University, Tempe, AZ 85287, USA

²School of Computing, Informatics, and Decision Systems Engineering, Arizona State University, Tempe, AZ 85281, USA

³Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109, USA

{hkerner, dfwellin, jim.bell, hbenamor}@asu.edu, kiri.l.wagstaff@jpl.caltech.edu

Abstract

In this work, we present a system based on convolutional autoencoders for detecting novel features in multispectral images. We introduce SAMMIE: Selections based on Autoencoder Modeling of Multispectral Image Expectations. Previous work using autoencoders employed the scalar reconstruction error to classify new images as novel or typical. We 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. We apply our methodology to the detection of novel geologic features in multispectral images of the Martian surface collected by the Mastcam imaging system on the Mars Science Laboratory *Curiosity* rover.

Introduction

The goal of novelty detection techniques is to identify new patterns in data that have not previously been observed (Markou and Singh 2003a; 2003b). Prior work has shown promising results in using autoencoder (Hinton and Salakhutdinov 2006) networks to detect novel observations and sensor readings. Autoencoders create mappings from high-dimensional input data, e.g., images, to a lower-dimensional manifold, along with a corresponding reverse mapping. Autoencoders thereby identify a compact representation that consists of the most salient features of the dataset.

An autoencoder can be trained in a purely unsupervised fashion to recognize novel input data that substantially deviate from previously seen data. The network is trained to reconstruct all training inputs with minimal loss. A new input sample can be fed into the trained network to yield a new reconstruction. The scalar reconstruction error can be used as a measure of the (un)familiarity of the autoencoder with this input. The assumption is that high reconstruction errors indicate *novel*, previously unseen input data. Such unsupervised training is particularly useful in applications like planetary exploration where the nature of novel features may not be well defined or known at all. Additionally, the cost of acquiring human labels for scientific data can be high due to limited access to domain experts.

One project that may greatly benefit from novelty detection is the Mars Science Laboratory *Curiosity* rover. *Curiosity* explores new regions on Mars every day as it is commanded to drive and make scientific observations in Gale crater. One instrument the rover uses to make geologic observations is the mast camera, or “Mastcam,” a pair of multispectral, color CCD imagers mounted on the rover’s mast about 2 meters above the surface (Bell et al. 2017; Malin et al. 2017; Grotzinger et al. 2012). Scientists only have a short time period after the data is received from the rover for planning follow-up observations (typically less than 12 hours), and this planning timeline will be even further reduced for future missions (for example, the goal is ≤ 5 hours for NASA’s Mars 2020 rover mission). For these reasons, there is a need for systems that can rapidly and intelligently extract information of interest from Mastcam and other science instrument data to focus on potential discoveries and avoid missed opportunities. Science planning and data analysis teams can benefit by spending their limited available time on the most promising, novel, or anomalous observations (e.g., Figure 1). Such a system must also provide explanatory visualizations that allow scientists to trust and understand how the system came to its conclusion – a need that is not supported by existing methods.

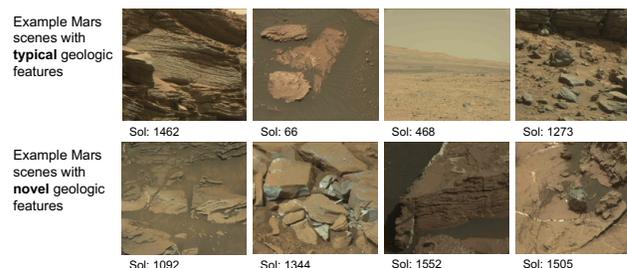


Figure 1: Example Mastcam images with typical (top) and novel (bottom) geologic features on Mars. Image credit: NASA/JPL/MSSS/ASU

In this work, we propose a neural network architecture called Selections based on Autoencoder Modeling of Multispectral Image Expectations (SAMMIE) for detecting and explaining novel features in multispectral images. We show how SAMMIE can identify novel geologic features in

Mastcam multispectral images of the Martian surface to facilitate scientific discovery and mission science planning. We provide a key insight that scalar reconstruction error does not always reliably indicate the novelty in this domain. As a solution, we argue and demonstrate that an autoencoder error map that captures spatial and spectral variation in errors yields a more suitable representation for estimating novelty in complex, high-dimensional data such as multispectral images. This method also enables human-comprehensible explanations and visualizations of detections.

Related Work

Novelty detection approaches. Most approaches to novelty detection focus on data from sensors or 3-channel color (RGB) images. Approaches for novelty detection in images typically detect anomalies in individual pixel spectra rather than spatial features within a multi-band image, (Kwon and Nasrabadi 2005; Bati et al. 2015). Traditional novelty detection includes statistical approaches such as density estimation, box-plots, and thresholds on distance from a class mean (Markou and Singh 2003a; Farrand, Merenyi, and Parente 2018), and neural network based approaches such as estimating output probability density functions, estimating prediction uncertainty, and thresholds on individual neuron activities (Markou and Singh 2003b).

Leveraging the success of deep neural networks at learning complex relationships in data, recent approaches use autoencoder networks for estimating novelty. In these approaches, an autoencoder is trained to minimize the reconstruction error for non-novel (typical) examples, typically the mean squared error between the input image and the image reconstructed by the autoencoder. Existing autoencoder-based approaches determine whether a new input is novel using a threshold or standard deviation metric on the scalar reconstruction error (Richter and Roy 2017; Japkowicz, Myers, and Gluck 1995; Xiong and Zuo 2016). While this has been sufficient for relatively simple datasets, it is not sufficient for all datasets as we show in this work.

Other deep learning approaches to novelty detection aim to quantify the novelty of inputs as a measure of confidence in the network’s output (Vasconcelos, Fairhurst, and Bisset 1995; Bishop 1994; Singh and Markou 2004; Chen et al. 2017). These are related to approaches for estimating the uncertainty of neural networks, for which methods based on Bayesian neural networks, ensembles of networks, and dropout have been proposed (Gal and Ghahramani 2016; Gal 2017; Mackay 1995; Lakshminarayanan, Pritzel, and Blundell 2017). Unsupervised novelty detection methods that learn representations for typical or non-novel datasets are also closely related to approaches to segmentation and objectness, where the model objective is to discover latent classes or features that are common in the data (Alexe, Deselaers, and Ferrari 2012; Badrinarayanan, Kendall, and Cipolla 2017).

Novelty detection in planetary exploration. Novelty detection for planetary science and astronomy datasets often employs Principal Component Analysis (PCA), a technique

used to identify patterns and outliers in data (Clegg et al. 2009; Dutta et al. 2007; Wagstaff et al. 2013). A limitation of linear methods such as PCA is that they can only represent simple (linear) relationships in the data. Linear methods can also be sensitive to potentially irrelevant variations in the data such as changes in position, orientation, and illumination of an object. For image applications, these methods also require transformation of the raw data (e.g., pixel values) into feature representations suitable for the learning algorithm (Wagstaff and Lee 2018).

Dataset

The Mastcam color imaging system on the Mars Science Laboratory (MSL) *Curiosity* rover acquires images for a variety of geologic and atmospheric studies (Bell et al. 2017). Each of Mastcam’s cameras, or “eyes,” has an eight-position filter wheel enabling image observations to be collected with “true color” (Bayer pattern broadband red, green, and blue) and with six narrow-band spectral filters spanning ~ 400 -1100 nm (visible and near-infrared) (Bell et al. 2017). The imagers have different focal lengths: 34 mm for the left eye and 100 mm for the right eye.

Our approach to novelty detection requires two datasets: one that represents the typical geology of Mars and one that contains expert-identified novel examples. To construct these datasets, we considered all Mastcam images acquired from sols (Martian days since landing) 3 to 1666 using all six narrow-band spectral filters by the left (M-34) and right (M-100) cameras. We use uncalibrated thumbnail versions of full-resolution multispectral images since these are the first products sent from the rover to Earth and thus are the earliest available tactical products. These images constitute a source dataset of 739 six-band images.

We identified novel examples in this dataset based on selections by an expert multispectral analyst on the Mastcam science team (Wellington et al. 2017a). We created a GUI for experts to draw 64×64 -pixel bounding boxes around areas of interest in the source dataset of 739 images (e.g., Figure 2). The multispectral image areas corresponding to these bounding boxes make up a dataset of 332 $64 \times 64 \times 6$ -pixel “novel” images. These images were excluded from the training set and later used for evaluation. After removing expert-selected novel images from the source dataset, we created a dataset of “typical” images by randomly cropping 64×64 -pixel frames from the source images for a total training dataset of 98,800 $64 \times 64 \times 6$ -pixel images. The source images are publicly available through the NASA Planetary Data System (PDS) and the labeled data is available at 10.5281/zenodo.1486196.

Methodology

We introduce a new methodology for detecting novel features in multispectral images as well as explaining *why* they were detected. The Selections based on Autoencoder Modeling of Multispectral Image Expectations (SAMMIE) system has three parts: a convolutional autoencoder for modeling the latent representations for typical data, a convolutional neural network classifier that classifies new examples

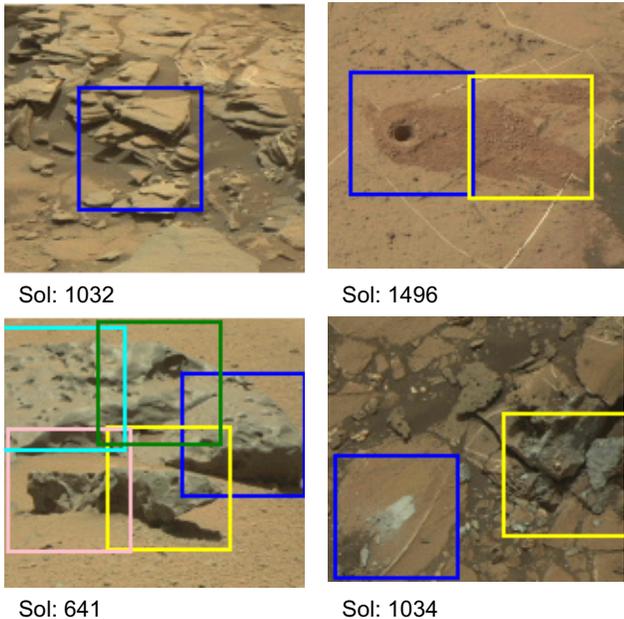


Figure 2: Example images selected by Mastcam multispectral experts for the novel image dataset. From top left to bottom right: an outcrop that may contain the mineral jarosite (Wellington et al. 2017b), a drill hole and tailings, the iron meteorite called “Lebanon” (Johnson et al. 2014), and a broken and wheel-scuffed surface. Image credit: NASA/JPL/MSSS/ASU

as containing novel content or not based on the autoencoder error map, and visualization tools for explaining the detection in a human-comprehensible way (see Figure 3).

Convolutional autoencoder to represent image content.

We use a convolutional autoencoder (CAE) architecture with three convolutional layers in the encoder and three transposed convolutional layers in the decoder. The layer architectures are $7 \times 7 \times 12$, $5 \times 5 \times 8$, and $3 \times 3 \times 3$ respectively for the encoder (reverse for the decoder). Since the input layer dimension is $64 \times 64 \times 6$ and the dimension of the bottleneck layer is $16 \times 16 \times 3$, examples are compressed by a factor of 32 before being reconstructed by the decoder. The loss function used for training is the mean squared error across all pixels and bands:

$$E(\mathbf{X}, \hat{\mathbf{X}}) = \frac{1}{NMK} \sum_{k=1}^K \sum_{j=1}^M \sum_{i=1}^N (x_{ij}^k - \hat{x}_{ij}^k)^2 \quad (1)$$

where x_{ij}^k and \hat{x}_{ij}^k are the pixel intensities at row i , column j , and band k in the input and reconstructed images respectively, N and M are the spatial dimensions of each image, and K is the number of bands. From the input and reconstructed images, we can construct a spatial-spectral error map $\delta(\mathbf{X}, \hat{\mathbf{X}})$, a $64 \times 64 \times 6$ tensor containing elements $(x_{ij}^k - \hat{x}_{ij}^k)^2$ for $i = 1, \dots, N$, $j = 1, \dots, M$, and $k = 1, \dots, K$.

Convolutional neural network for novelty detection.

The CAE error map identifies novel features at the pixel level but does not estimate the novelty of the image as a whole. Previous approaches used a threshold on the scalar reconstruction error (Eqn. 1) to detect novelty. In our dataset, the scale of novel features is highly variable, and typical and novel error values are not linearly separable (Figure 4). We therefore use a convolutional neural network (CNN) to predict whether the CAE error map represents truly novel features or not. The CNN is a binary classifier with two convolutional layers with sizes $5 \times 5 \times 32$ and $5 \times 5 \times 64$, followed by one dense layer with 512 units, a dropout layer, and a weighted softmax layer.

Explanations to visualize novel content.

Comparing the input image with the image reconstructed by the autoencoder reveals which features were considered novel by observing what was “lost” in the reconstruction (Figure 3). The explanation component of SAMMIE builds on this intuition by using the CAE error map to highlight the novel content in an image for the end user. This enables a richer explanation of the detection through pixel-wise comparisons of the observed spectrum (reflectance in one pixel location across all six bands) to the spectrum SAMMIE expected to find in highlighted novel regions. We explore explanations in more detail in the Explanations section.

Experiments

We trained the convolutional autoencoder (CAE) component of SAMMIE on the typical images, then used the resulting error output from the CAE to train five different novelty classifiers. We compare these results with direct classification of the multispectral input image to assess the contribution of the autoencoder step.

Classification of Autoencoder Error Output

We compared the novelty detection performance of SAMMIE using the same convolutional autoencoder combined with different novelty classifiers: Naive Bayes, a feed-forward neural network (FFNN), a CNN, and Inception-V3 pre-trained on the ImageNet database (Szegedy et al. 2015; Deng et al. 2009). We trained the CNN and Inception-V3 classifiers using CAE error maps for 98,700 typical (negative) examples and 300 positive examples from the novel dataset described in the Dataset section. We trained the Naive Bayes and FFNN classifiers using only the mean of each CAE error map for the same examples. To correct for the extreme class imbalance, we assigned a weight to positive examples that is inversely proportional to the positive class occurrence for all experiments except those using Inception-V3: $w_i = 1 + c_i \left(\frac{98,700}{300} - 1 \right)$ for example i having class $c_i \in \{0, 1\}$. We used a fixed test set of 132 (100 negative and 32 positive) randomly selected examples for all experiments. We describe each experiment below and compare their performance in Table 1. Figure 5 compares the receiver operating characteristics (ROC) and area under the curve (AUC) computed from the ROC for each classifier. The accuracies reported in Table 1 were computed using the

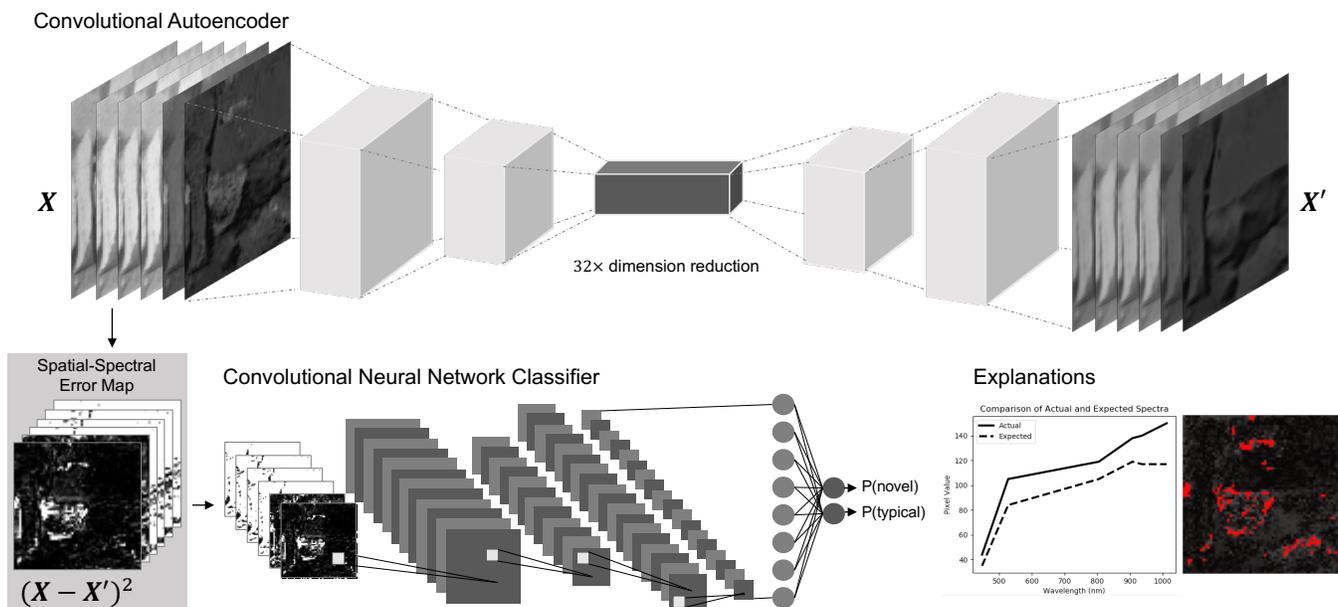


Figure 3: SAMMIE system comprising a convolutional autoencoder with matrix-typed reconstruction error, a convolutional neural network for classification, and products for explaining detections. Image credit: NASA/JPL/MSSS/ASU

threshold on posterior probability from the ROC curve that maximized classification accuracy on the test set.

Naive Bayes. Previous work (Richter and Roy 2017; Japkowicz, Myers, and Gluck 1995; Xiong and Zuo 2016) has shown that a threshold on the autoencoder scalar reconstruction error (the mean squared error between the reconstructed and input examples) is sufficient for discriminating between typical and novel inputs. We trained a Naive Bayes classifier to predict novelty using the scalar error value (Figure 4). It is clear from Figure 4 that there is significant overlap between the distributions of mean squared error in the typical image and novel image datasets. Maximum accuracy of 78.0% was achieved with the posterior threshold 0.415.

Feed-forward neural network. As an alternative to the Naive Bayes classifier, we trained a feed-forward neural network classifier to predict whether an example was typical or novel based on the scalar reconstruction error. The classifier consisted of three dense layers of size 5 – 10 – 5. Maximum accuracy of 78.0% was achieved with the threshold 0.627.

Inception V3 network. Since the Inception network requires 3-channel images as input, we created two separate input datasets for the shorter wavelength Mastcam filters (447, 527, and 805 nm for the M-100; 445, 527, and 676 nm for the M-34) and the longer wavelength filters (908, 937, and 1013 nm for the M-100; 751, 867, 1012 nm for the M-34). We refer to these networks in Table 1 as Inception-V3 (short) and Inception-V3 (long). We fine-tuned the final layer of Inception-V3 using TensorFlow (Abadi et al. 2015). We found that this model achieved better performance when

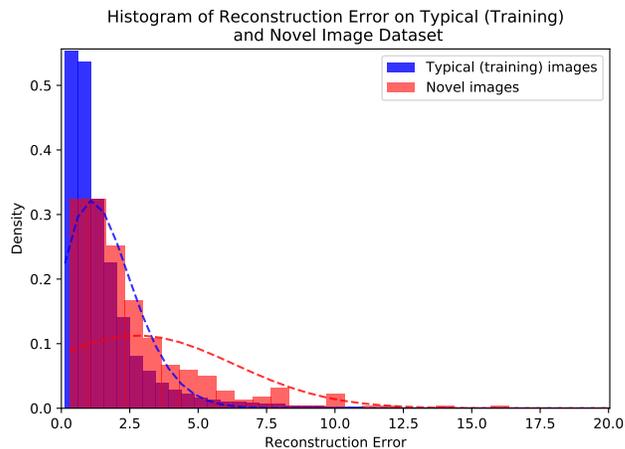


Figure 4: Histogram of autoencoder reconstruction error (mean squared error) for typical and novel images showing significant overlap between the two distributions.

examples were not weighted during training to correct for the class imbalance as in the other classifiers. Maximum accuracy of 84.8% was achieved with the threshold 0.217 for Inception-V3 (short) and 79.5% with the threshold 0.594 for Inception-V3 (long).

SAMMIE. SAMMIE uses a convolutional neural network to classify the CAE error map as described in the Dataset section. The input to the CNN was the $64 \times 64 \times 6$ error map. Maximum accuracy of 96.2% was achieved with the

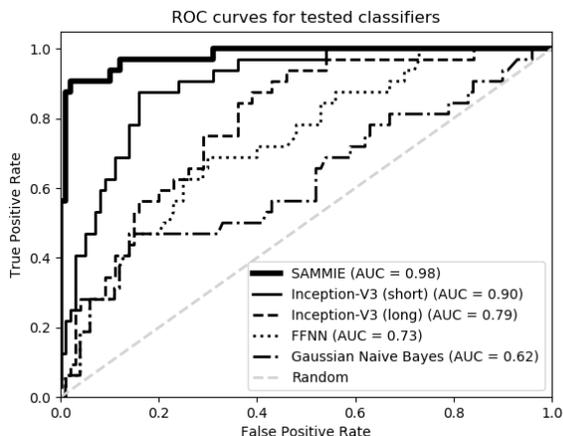


Figure 5: Receiver operating characteristics (ROC) curves for five compared classifiers of autoencoder error output.

threshold 0.802.

Direct Classification

In a second experiment (Table 2), we tested the ability of three models to classify the novelty of the multispectral input image directly, rather than passing the input first through the CAE. As in the previous experiment, we fine-tuned separate Inception-V3 models using shorter and longer wavelength Mastcam filters. We used the classification threshold that yielded the highest accuracy as in the previous experiment. We found that the Inception networks performed better when classifying the original input directly than when using the autoencoder as a pre-processor. This is likely because the features represented in the original input are more similar to the natural images the Inception network was trained on compared to the images of reconstruction error that the autoencoder produces. Although the performance of the Inception-V3 (short) model comes close to SAMMIE, it does not provide a means for visualizing intuitive explanations and would not be expected to perform well given features not represented in the training set.

Qualitative results

Figure 6 shows images from the test set that were classified as typical or as novel with at least 99.9% probability. All were correctly classified. More than half of the typical example images contain rover hardware, such as the calibration target (dark vertical pole with sphere on top). This target is imaged very frequently and thus is easily recognized by SAMMIE. Other images in the typical set do not appear to contain significant spectral diversity compared with the images in the novel set. Of the images in the novel set, two contain meteorites, and several include marks left by the rover brushing or drilling the surface.

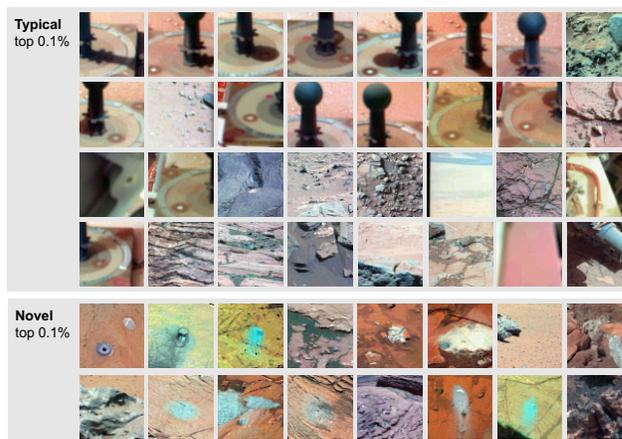


Figure 6: Test set images classified as typical or novel with probability $\geq 99.9\%$. Image credit: NASA/JPL/MSSS/ASU (images are stretched to emphasize spectral diversity).

Explanations

When analyzing multispectral images, scientists typically use a spectral analysis tool to inspect the spectrum in a region of interest within the image (single or groups of pixels). They compare the observed spectra to known spectral patterns and characteristics for different materials to come up with interpretations for the observed data (Wellington et al. 2017a).

From the error map produced by SAMMIE, we can visualize for any pixel in the image the detected novel content alongside the observed 6-filter pixel spectrum and the spectrum that SAMMIE reconstructed (Figure 7). In the image of an iron meteorite in Figure 7 (top), SAMMIE identified the meteorite in the image (region 1) as novel. The explanation shows that the rock has a higher signal in filter 6 (1013 nm) than expected, which is consistent with the increase in near-infrared reflectance values between the wavelengths of filter 5 (937 nm) and 6 that is typical of iron meteorites, relative to native Martian materials (Gaffey 1976; Wellington et al. 2017a). In contrast, in a region of the image that SAMMIE identified as typical (region 2), the explanation shows that the spectrum SAMMIE expected matches well with the actual spectrum in that region. The bottom image of Figure 7 shows a raised ridge that was crushed by the rover’s wheel. Comparing the spectra in Figure 7 (bottom) shows that the region 1 signal in the near-infrared wavelengths starting around 750 nm is much lower relative to filter 2 (527 nm), creating a downturn between the signal in the near-infrared filters that is consistent with the decrease in near-infrared reflectance values observed for magnesium-rich ridges by other instruments on MSL (Johnson et al. 2015; Leveille et al. 2014).

Conclusions and Future Work

We presented a system based on convolutional autoencoders for novelty detection in multispectral images. Our system, Selections based on Autoencoder Modeling of Multispec-

| Model | Accuracy | True Pos. | False Pos. | True Neg. | False Neg. | AUC |
|----------------------|--------------|-----------|------------|-----------|------------|-------------|
| Naive Bayes | 78.0% | 9 | 6 | 94 | 23 | 0.62 |
| FFNN | 78.0% | 9 | 6 | 94 | 23 | 0.73 |
| Inception-V3 (long) | 79.5% | 8 | 3 | 97 | 24 | 0.79 |
| Inception-V3 (short) | 84.8% | 28 | 16 | 84 | 4 | 0.90 |
| SAMMIE | 96.2% | 28 | 1 | 99 | 4 | 0.98 |

Table 1: Test set performance for image novelty classification based on CAE error output.

| Model | Accuracy | True Pos. | False Pos. | True Neg. | False Neg. | AUC |
|----------------------|----------|-----------|------------|-----------|------------|------|
| CNN | 82.6% | 13 | 4 | 96 | 19 | 0.87 |
| Inception-V3 (long) | 87.1% | 16 | 1 | 99 | 16 | 0.91 |
| Inception-V3 (short) | 90.1% | 21 | 2 | 98 | 11 | 0.95 |

Table 2: Test set performance for direct novelty prediction from multispectral input images.

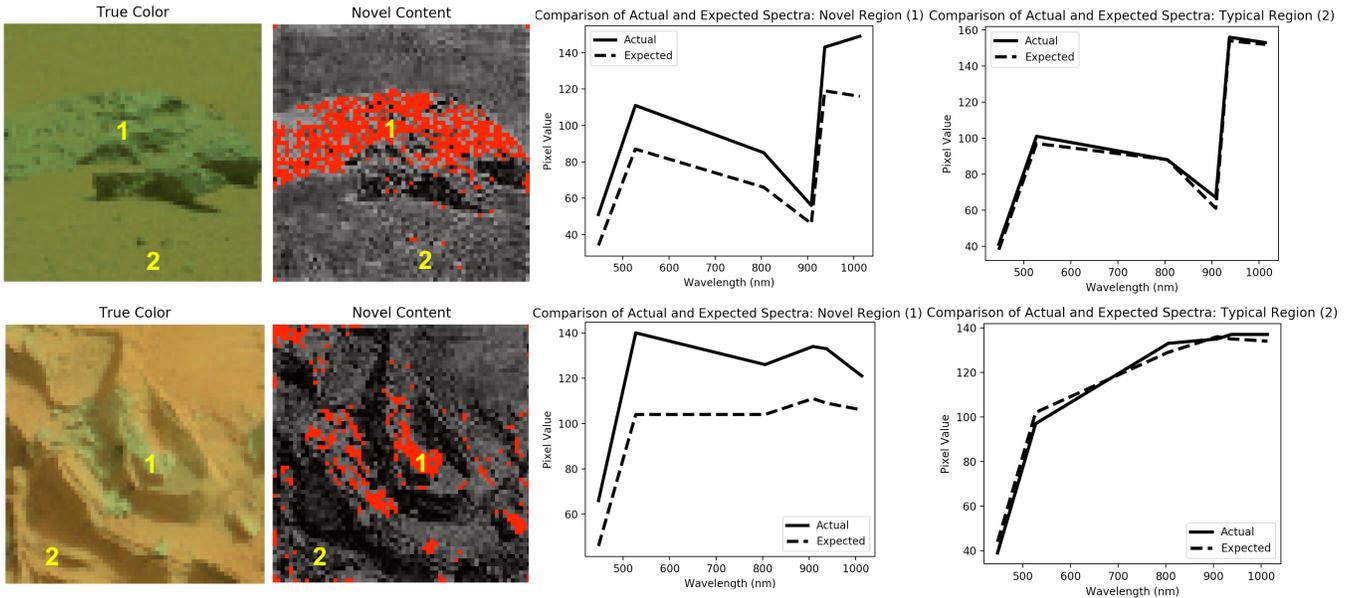


Figure 7: Example explanations for two images identified by SAMMIE as novel. Areas highlighted in red have the largest reconstruction error (novelty). Top row: an iron meteorite. Bottom row: a raised, fracture-filling ridge crushed by the rover’s wheel. Comparing the actual and expected pixel values in each band helps multispectral image analysts determine why the indicated feature might be novel. Image credit: NASA/JPL/MSSS/ASU (sol 346, sequence mcam01405; sol 164, mcam00883).

tral Image Expectations (SAMMIE), uses autoencoder error maps to make classifications of images as novel or typical with respect to the training set. This enables our system to be sensitive to subtle novel features within an image as well as robust to potentially false indicators of spectral novelty, such as overexposure or saturation. In addition, our model enables human-comprehensible explanations of detections. We showed how this novelty detection method could be applied to multispectral images from the Mars Science Laboratory *Curiosity* Mastcam instrument investigation.

In future work, we will continue developing the explanation capability of SAMMIE and conduct experiments to measure the benefits of using SAMMIE in practice for MSL science operations. We will explore the ability of autoencoder representations to transfer between similar instru-

ments and images of similar scenes, such as between Mastcam and the Pancam imaging system on the Mars Exploration Rovers *Spirit* and *Opportunity* (Bell III et al. 2008). Additionally, we plan to explore the use of autoencoders for change detection and the novelty detection capability of other generative models including Generative Adversarial Networks (GANs) (Goodfellow et al. 2014).

Acknowledgments

This work was funded by NASA STTR #80NSSC17C0035 and NASA/JPL funding from the Mars Science Laboratory Mastcam instrument investigation. It was carried out in part at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

References

- Abadi, M.; Agarwal, A.; Barham, P.; Brevdo, E.; Chen, Z.; Citro, C.; Corrado, G. S.; Davis, A.; Dean, J.; Devin, M.; Ghemawat, S.; Goodfellow, I.; Harp, A.; Irving, G.; Isard, M.; Jia, Y.; Jozefowicz, R.; Kaiser, L.; Kudlur, M.; Levenberg, J.; Mané, D.; Monga, R.; Moore, S.; Murray, D.; Olah, C.; Schuster, M.; Shlens, J.; Steiner, B.; Sutskever, I.; Talwar, K.; Tucker, P.; Vanhoucke, V.; Vasudevan, V.; Viégas, F.; Vinyals, O.; Warden, P.; Wattenberg, M.; Wicke, M.; Yu, Y.; and Zheng, X. 2015. TensorFlow: Large-scale machine learning on heterogeneous systems. Software available from tensorflow.org.
- Alexe, B.; Deselaers, T.; and Ferrari, V. 2012. Measuring the Objectness of Image Windows. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 34(11):2189–2202.
- Badrinarayanan, V.; Kendall, A.; and Cipolla, R. 2017. SegNet: A Deep Convolutional Encoder-Decoder Architecture for Image Segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 39(12):2481–2495.
- Bati, E.; Çalkan, A.; Koz, A.; and Alatan, A. A. 2015. Hyperspectral Anomaly Detection Method based on Auto-encoder. In Bruzzone, L., ed., *Proceedings of the SPIE, Image and Signal Processing for Remote Sensing*, volume 9643, 96430N.
- Bell, J. F.; Godber, A.; McNair, S.; Caplinger, M. A.; Maki, J. N.; Lemmon, M. T.; Van Beek, J.; Malin, M. C.; Wellington, D.; Kinch, K. M.; Madsen, M. B.; Hardgrove, C.; Ravine, M. A.; Jensen, E.; Harker, D.; Anderson, R. B.; Herkenhoff, K. E.; Morris, R. V.; Cisneros, E.; and Deen, R. G. 2017. The Mars Science Laboratory Curiosity Rover Mastcam Instruments: Preflight and In-flight Calibration, Validation, and Data Archiving. *Earth and Space Science* 4(7):396–452.
- Bell III, J.; Calvin, W.; Farrand, W.; Greeley, R.; Johnson, J.; Joliff, B.; Morris, R.; Sullivan, R.; Thompson, S.; Wang, A.; Weitz, C.; and Squyres, S. 2008. *Mars Exploration Rover Pancam Multispectral Imaging of Rocks, Soils, and Dust in Gusev Crater and Meridiani Planum*. Cambridge University Press.
- Bishop, C. 1994. Novelty Detection and Neural Network Validation. In *IEE Proceedings - Vision, Image, and Signal Processing*, volume 141, 217–222.
- Chen, J.; Sathe, S.; Aggarwal, C.; and Turaga, D. 2017. Outlier Detection with Autoencoder Ensembles. In *Proc. of the SIAM International Conference on Data Mining*, 90–98.
- Clegg, S. M.; Sklute, E.; Dyar, M. D.; Barefield, J. E.; and Wiens, R. C. 2009. Multivariate Analysis of Remote Laser-Induced Breakdown Spectroscopy Spectra using Partial Least Squares, Principal Component Analysis, and Related Techniques. *Spectrochimica Acta Part B: Atomic Spectroscopy* 64(1):79–88.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Kai Li; and Li Fei-Fei. 2009. ImageNet: A Large-Scale Hierarchical Image Database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition (CVPR'09)*, 248–255. IEEE.
- Dutta, H.; Giannella, C.; Borne, K.; and Kargupta, H. 2007. Distributed Top-K Outlier Detection from Astronomy Catalogs using the DEMAC System. In *Proceedings of the 2007 SIAM International Conference on Data Mining*, 473–478.
- Farrand, W.; Merenyi, E.; and Parente, M. 2018. *Hyper- and Multispectral VNIR Imaging Analysis*. Cambridge Univ. Press, in press.
- Gaffey, M. J. 1976. Spectral Reflectance Characteristics of the Meteorite Classes. *Journal of Geophysical Research* 81(5):905–920.
- Gal, Y., and Ghahramani, Z. 2016. Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning. In *Proceedings of the 33rd International Conference on Machine Learning (ICML)*, volume 3, 1651–1660.
- Gal, Y. 2017. *Uncertainty in Deep Learning*. Ph.D. Dissertation, University of Cambridge.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative Adversarial Nets. In *Advances in Neural Information Processing Systems (NIPS) 27*, 2672–2680.
- 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 Reviews* 170(1-4):5–56.
- Hinton, G. E., and Salakhutdinov, R. R. 2006. Reducing the Dimensionality of Data with Neural Networks. *Science* 313(5786):504–507.
- Japkowicz, N.; Myers, C.; and Gluck, M. 1995. A Novelty Detection Approach to Classification. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence (IJCAI)*, volume 1, 518–523.
- Johnson, J. R.; Bell, J. F., I.; Gasnault, O.; Le Mouelic, S.; Rapin, W.; Bridges, J.; and Wellington, D. F. 2014. First Iron Meteorites Observed By the Mars Science Laboratory (MSL) Rover Curiosity. In *Proceedings of the American Geophysical Union Fall Meeting*.
- Johnson, J. R.; Bell III, J.; Bender, S.; Blaney, D.; Cloutis, E.; DeFlores, L.; Ehlmann, B.; Gasnault, O.; Gondet, B.; Kinch, K.; Lemmon, M.; Le Mouelic, S.; Maurice, S.; Rice, M.; and Wiens, R. 2015. ChemCam Passive Reflectance Spectroscopy of Surface Materials at the Curiosity Landing Site, Mars. *Icarus* 249:74–92.
- Kwon, H., and Nasrabadi, N. 2005. Kernel RX-algorithm: a Nonlinear Anomaly Detector for Hyperspectral Imagery. *IEEE Transactions on Geoscience and Remote Sensing* 43(2):388–397.
- Lakshminarayanan, B.; Pritzel, A.; and Blundell, C. 2017. Simple and Scalable Predictive Uncertainty Estimation using Deep Ensembles. In *Advances in Neural Information Processing Systems (NIPS) 30*.
- Leveille, R. J.; Bridges, J.; Wiens, R. C.; Mangold, N.; Cousin, A.; Lanza, N.; Forni, O.; Ollila, A.; Grotzinger,

- J.; Clegg, S.; Siebach, K.; Berger, G.; Clark, B.; Fabre, C.; Anderson, R.; Gasnault, O.; Blaney, D.; Deflores, L.; Leshin, L.; Maurice, S.; and Newsom, H. 2014. Chemistry of Fracture-filling Raised Ridges in Yellowknife Bay, Gale Crater: Window into Past Aqueous Activity and Habitability on Mars. *Journal of Geophysical Research: Planets* 119(11):2398–2415.
- Mackay, D. J. C. 1995. Probable Networks and Plausible Predictions – a Review of Practical Bayesian Methods for Supervised Neural Networks. *Network: Computation in Neural Systems* 6(3):469–505.
- Malin, M. C.; Ravine, M. A.; Caplinger, M. A.; Tony Ghaemi, F.; Schaffner, J. A.; Maki, J. N.; Bell, J. F.; Cameron, J. F.; Dietrich, W. E.; Edgett, K. S.; Edwards, L. J.; Garvin, J. B.; Hallet, B.; Herkenhoff, K. E.; Heydari, E.; Kah, L. C.; Lemmon, M. T.; Minitti, M. E.; Olson, T. S.; Parker, T. J.; Rowland, S. K.; Schieber, J.; Sletten, R.; Sullivan, R. J.; Sumner, D. Y.; Aileen Yingst, R.; Duston, B. M.; McNair, S.; and Jensen, E. H. 2017. The Mars Science Laboratory (MSL) Mast Cameras and Descent Imager: Investigation and Instrument Descriptions. *Earth and Space Science* 4(8):506–539.
- Markou, M., and Singh, S. 2003a. Novelty Detection: a Review - Part 2: Neural Network Based Approaches. *Signal Processing* 83(12):2499–2521.
- Markou, M., and Singh, S. 2003b. Novelty Detection: a Review-Part 1: Statistical Approaches. *Signal Processing* 83(12):2481–2497.
- Richter, C., and Roy, N. 2017. Safe Visual Navigation via Deep Learning and Novelty Detection. In *Proceedings of Robotics Science and Systems (RSS)*.
- Singh, S., and Markou, M. 2004. An Approach to Novelty Detection Applied to the Classification of Image Regions. *IEEE Transactions on Knowledge and Data Engineering* 16(4):396–406.
- Szegedy, C.; Liu, W.; Jia, Y.; Sermanet, P.; Reed, S.; Anguelov, D.; Erhan, D.; Vanhoucke, V.; and Rabinovich, A. 2015. Going Deeper With Convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Vasconcelos, G.; Fairhurst, M.; and Bisset, D. 1995. Investigating Feedforward Neural Networks with Respect to the Rejection of Spurious Patterns. *Pattern Recognition Letters* 16(2):207–212.
- Wagstaff, K. L., and Lee, J. 2018. Interpretable discovery in large image data sets. In *Proc. of the Workshop on Human Interpretability in Machine Learning (WHI)*, 107–113.
- Wagstaff, K. L.; Lanza, N. L.; Thompson, D. R.; Dietterich, T. G.; and Gilmore, M. S. 2013. Guiding Scientific Discovery with Explanations Using DEMUD. In *Twenty-Seventh AAAI Conference on Artificial Intelligence*, 905–911.
- Wellington, D. F.; Bell III, J. F.; Johnson, J. R.; Rice, M. S.; Fraeman, A. A.; and Horgan, B. 2017a. VIS/NIR Spectral Differences of Materials within Gale Crater, Mars: Parameterization of MSL/Mastcam Multispectral Observations. In *Proc. of the 48th Lunar and Planetary Science Conference*.
- Wellington, D. F.; Bell, J. F.; Johnson, J. R.; Kinch, K. M.; Rice, M. S.; Godber, A.; Ehlmann, B. L.; Fraeman, A. A.; and Hardgrove, C. 2017b. Visible to Near-Infrared MSL/Mastcam Multispectral Imaging: Initial Results from Select High-Interest Science Targets within Gale Crater, Mars. *American Mineralogist* 102(6):1202–1217.
- Xiong, Y., and Zuo, R. 2016. Recognition of Geochemical Anomalies using a Deep Autoencoder Network. *Computers and Geosciences* 86:75–82.