Autonomous Mapping of Outcrops Using Multiclass Linear Discriminant Analysis

R. Francis*, K. McIsaac*, D. R. Thompson**, G. R. Osinski*

*Centre for Planetary Science and Exploration, University of Western Ontario
contact: raymond.francis@cpsx.uwo.ca

**Jet Propulsion Laboratory, California Institute of Technology

Abstract

The growing capability of robotic planetary exploration missions brings increasingly constrained data budgets. The autonomous prioritization, processing, and even acquisition of science data could allow more scientific information and discovery per megabit of transmission – increasing the scientific value and the pace of operations.

Here we explore the problem of analyzing rock outcrops using only colour photography. Elementary image processing techniques generate a feature space incorporating colour and visual texture information from an image. Then, a machine learning technique, multiclass linear discriminant analysis (MDA), is used to learn the visual components that distinguish adjacent rock types from each other, and a vector clustering technique is used to segment images of the same and similar outcrops. The approach is tested successfully in a variety of geological settings representative of those expected in planetary surface exploration, and shows promise for extension to other imaging modalities and processing techniques.

1 Introduction

Planetary exploration missions have deployed progressivley more capable and complex platforms to explore the bodies in the solar system. The Mars Science Laboratory rover, for example, carries an extensive suite of remote-sensing, contact, and internal instruments [1], together offering far greater capacity to generate scientific data than previous missions. Communications capabilities continue to be an important limitation for these missions, however, restricting the amount of data that can be returned to Earth. This has an additonal consequence of restricting the speed of operations, as many tasks – such as approaching a rock outcrop and placing an instrument against it – require several human-in-the-loop steps in decision-making. Each such step requires the transmission of data to the Earth, its inspection, and the transmission of resultant commands to the spacecraft.

1.1 Improving scientific throughput

These missions could realize a greater throughput of science information by achieving these tasks more quickly, or by devoting a larger fraction of the downlink to science-relevant images as opposed to intermediate images intended for human decision-making. Achieving either – more science per time, or more science per kilobit – may be possible by allowing the robotic system to make decisions on its own.

One tool for aiding such decision-making by the robot is autonomous interpretation of images. A major task of operations scientists is finding interesting features in downlinked images so that they may be targeted for further investigation. New software has already begun to allow a degree of autonomous interpretation and improved scientific return, as in the case of the AEGIS software developed for the Mars Exploration Rovers [2], which uses on board image processing to prioritize images for downlink, based on the detection of rocks sitting on the plain surrounding the rover.

1.2 Geological investigations: In-place materials and contacts

For a surface mission with a geological focus, an example of an interesting feature is an outcrop in which visual inspection suggests the presence of more than one type of rock. In terrestrial field geology, these features are called *contacts* between *geological units*, and while they have a variety of forms and origins, they are important sources of information about the nature and history of the rocks in a region. The information they carry is particularly rich if the materials are still *in-place*, that is, remaining in the position and context of their formation. This makes outcrops valuable sites of investigation, over loose material that cannot be clearly connected to its place of origin.

Impact craters, ubiquitous on solid planetary surfaces and of interest for both geological and astrobiological investigations [3], are formed by violent and turbulent processes that result in numerous emplacements of mixed rock types, such as impact breccias, melt rocks, and ejecta deposits [4]. This chaotic process also tends to produce

features of vertical relief, often including exposed outcrops of rock, particularly in environments where the rate of erosion is low. Finding such outcrops, particularly those displaying geological contacts, is critical in investigations of such widely-varying settings as impact craters, sedimentary environments, and volcanic flows.

1.3 Present work

The present work aims to develop image processing tools which allow the segmentation of images of rock outcrops along geological units, as a step towards producing a capability for detection and mapping of outcrops displaying geological contacts. We formalize this problem as one of unsupervised image segmentation into geologic surface types. We test the hypothesis that Mahalanobis metric learning, trained by exemplar scenes, can improve the fidelity of such segmentations to an expert interpretation. The experiments show that unsupervised image segmentation produces geologically relevant categories from simple colour and texture primitives when analyzed with appropriate distance metrics. Section 2 describes the history of similar work in the exploration context. The design of the algorithm is presented in section 3, and the design of an experiment to test it is described in section 4. The results of tests on field data in a variety of geologic settings are presented and discussed in section 5.

2 Existing techniques

2.1 Geological classification

Geology on Earth relies site visits by a trained analyst, but planetary exploration programs have researched automated techniques. Previous efforts have identified specific features in outcrops, especially sedimentary layering [5] [6]. Further efforts saw significant successes in detecting loose rocks resting on the ground [7] [8], by a variety of techniques with varying results [9]. Such systems eventually became capable enough to be used to guide robotic decision making [10], and in recent years to be deployed on the Mars Exploration Rover Opportunity to improve the return of scientific data [11]. But with highly variable visual appearance, gradational boundaries, and complex boundary morphology, separation of geological materials within an outcrop is even more challenging a task than recognizing rocks against a background. At least one current project is working to apply information from image texture to the problem [12] but to our knowledge the present work is the first attempt to identify geologic contacts within a single image in a wholly unsupervised fashion.

2.2 Distance metric learning

In the presence of correlations and noise dimensions in the input space, it can be quite difficult to find a representation where the classes of interest naturally separate from each other. Metric Learning seeks a distance metric, or equivalently, a transformation of the input space, to maximize task performance. These methods typically rely on a training set of distinct classes from the problem domain. They optimize the distance metric to maximize the distance between the dissimilar classes. This new representation reflects semantic distinctions of interest so that a wholly unsupervised algorithm can recover them.

Many such algorithms involve fitting a Mahalanobis metric, expressed as a simple linear projection that we will write here as a matrix A. Following the notation of [13], applying A to each sample pair (x_i, x_j) produces a Mahalanobis distance in the original sample space:

$$M(x_i, x_j) = (A^T x_i - A^T x_j)^T (A^T x_i - A^T x_j)^T \\ = (x_i - x_j) A A^T (x_i - x_j) \tag{1}$$

Note that the matrix $M = AA^T$ is symmetric, positive semi-definite. There are many ways to find this matrix. The most common approach, Multiclass Discriminant Analysis, is an extension of classical Linear Discriminant Analysis; for k distinct classes, it forms A by the eigenvectors associated with the top k-1 eigenvalues of $M_w^{-1} M_b$ Here M_w is the within class scatter matrix and M_b is the between class scatter matrix. The resulting transformation minimizes the determinant of the former and maximizes the determinant of the latter. Other linear distance metric learning algorithms include Information Theoretic Metric Learning [14], Locally Discriminative Gaussians [15], and NCA [16]. A comparison of all such techniques is beyond the scope of this paper, but LDA-based methods often perform comparably to iterative approaches and basic MDA is sufficient to evaluate our hypothesis.

3 Method

3.1 Strategy

The approach used in this work begins from the observation that rocks which are visually distinct from each other are often distinct in several ways - by colour, by the orientation of linear features in the rock surface ('fabric', in the geological sense), or visual texture, such as from layering, weathering, fracturing, or grain size. Very often separate geological units are visually distinguishable from each other in more than one of these characteristics. We attempt to exploit this property by applying a technique that finds groups of pixels which vary together in several visual features.

3.2 Channel set

We begin by processing the image to produce several data products relating to colour, texture, and other

visual attributes at each pixel. Each such data product is called hereafter a visual '*channel*', and represents an array of values corresponding to each pixel in the input image. For initial tests, a basic feature set was used consisting of seven channels:

1. The grayscale representation of the colour image;

2. The red channel of the colour image;

3. The green channel of the colour image;

4. The blue channel of the colour image;

5. The ratio, at each pixel, of the blue and red channels;

6. The difference, at each pixel, of the blue and red channels;

7. A brightness map produced by first taking the magnitude of the image gradient at each pixel, then passing a kernel over the result which sums the values of all pixels in a small radius. This channel is intended to respond to the local density of edge features in regions of the image.

An extended feature set uses these same channels along with eight more provided by the MR8 filter bank [17], in an effort to further emphasize textural information. In principle, many more channels can be designed and included, but this present work reports only on results using the above basic and extended feature sets.

For either implementation, the data produced in creating the *n*-dimensional feature space is represented as a set of *n*-dimensional vectors, with each pixel represented by a vector composed of that pixel's corresponding values from each visual channel.

3.3   Learning step

We first train MDA by using a dataset of labeled images from the same locale. This is relevant for spacecraft operations where a rover is travelling tens or hundreds of metres per command cycle, and geologic surface types will be somewhat similar to categories that have already been seen in previous images. Scientists could train such a system on the ground and then transmit the compact transformation matrix to the rover, enabling it to recognize appropriate features in new images. For each set of training data, we formed 2 to 3 classes from the categories of interest, and learned an MDA representation based on this training data. MDA permits solutions with a rank up to k-1 where k is the number of classes. We then applied the low-rank transform to other images from the locale not used during training, producing an unbiased estimate of task performance on a new scene.

To effect the segmentation, the feature space vectors are transformed to the MDA-learned representation, then clustered by proximity in the *n*-dimensional feature space. As a baseline, the k-means clustering technique is used, with other clustering techniques possible.

3.4   Assessing the segmentation

Candidate segmentations are assessed by comparison to manually-labeled reference segmentations using the Adjusted Rand Index [18]. The Rand Index is a figure of merit that counts the number of pairs of pixels which are, in both segmentations, assigned to the same segment, as a fraction of the total number of pixels. Normalized against random chance, it becomes the Adjusted Rand Index (ARI), which has a value of zero for a pixel segmentation performing the same as random assignment, and a value of one for a segmentation which is identical to the reference.

4   Experiment design

The technique was tested on imagery from a variety of geological settings, including several types of volcanic deposits in Mars-analogue sites in the Mojave desert, California; impact breccias from the Sudbury impact crater in Ontario, Canada; and a clay-rich sedimentary setting with visible calcium sulfate veins in Gale Crater, Mars. In each geological setting, the system is trained using a representative image showing the characteristic local rock types. The trained system then segments both the training image, and new images from the same locality. Three cases are tested in each locality, each using a different feature space:

- The basic feature space (described in section 3.2), without applying the learned vector ("No learning")

- The basic feature space, with the learned transformation applied

- The extended feature space including the MR8 filter bank, with MDA learning on that larger space

In each case, the segmentation is compared to a reference, human-labeled segmentation, using the Adjusted Rand Index.

For the locales on Earth, the photographs used were captured with a handheld digital SLR camera. Lighting conditions varied from full insolation to full shade; cast shadows that covered only a portion of the scene were avoided, though shadows created by in-scene relief are unavoidably present in several cases. The photographs were selected to show a variety of geological materials with visible contacts, and a variety of contact types (sharp, gradational, highly complex) and morphologies (adjacent massive units, layered materials, clasts within a matrix). Artifical objects often included in geological imaging, such as hammers and rulers to provide a reference for scale, were

Table 1. Segmentation algorithm performance. Values are the Adjusted Rand Index, as described in section 3.4

Scene type, *Number of classes*	Image reference	No learning	Learning on basic feature space	Learning on extended feature space
A	*1277*	0.814	0.968	0.987
2	1261	0.912	0.922	0.936
	1281	0.912	0.961	0.973
B	*0272*	0.562	0.759	0.816
2	0743	0.378	0.657	0.715
	0283	0.348	0.538	0.609
C	*0199*	0.565	0.943	0.950
3	0200	0.372	0.795	0.749
	0230	0.245	0.587	0.563
D	*0495*	0.323	0.693	0.742
3	0497	0.289	0.616	0.633
	0511	0.350	0.589	0.665
E	*9726*	0.774	0.920	0.946
2	9735	-0.072	0.790	0.793
	9745	0.761	0.837	0.868
F	*s133r1*	0.084	0.769	0.828
2	s133ls2	0.053	0.220	0.718
	s133ls3	-0.002	0.018	0.313

Figure 1. Example of scene type A. Basalt blocks and sand, image number 1281, and its segmentation result.

Figure 2. Example of scene type B. Massive basalt and lahar deposit, image number 0272, and its segmentation result.

intentionally excluded. These procedures allowed the experiment to employ a set of fully natural scenes with no intrusions, but somewhat optimized shadow for the local topography. Such conditions are representative of those to be expected in planetary surface imaging with a robotic platform.

Images from Mars were obtained by the left and right Mastcam imagers of the Mars Science Laboratory rover in the context of the mission science team's investigation of the Yellowknife Bay locality of Gale Crater. Scenes of the desired lithologies having broadly similar dust cover were selected from the area visited by the rover on sol 133 of the mission, with views of the rover hardware excluded.

5 Results

We tested the algorithm on a variety of scenes having diverse rock types and boundary shapes. These scene types are described in the following section, and include a variety of volcanic, impact, and sedimentary settings, each showing clear contacts between distinct geological units.

The results of the segmentation are shown in Table 1. For each scene type, one image of a representative scene

was used to train the algorithm. The feature space transformation learned using this image was used in segmenting this same image, and two more images of the same type of geological materials, taken under similar lighting conditions at an adjacent site, with the image fields of view not overlapping that of the training image. For each scene type, the image used for training is that marked by by an italicised reference number in Table 1.

5.1 Scene types

5.1.1 Type A: Basalt blocks and sand

A mixture of irregular, vesiculated basalt blocks, surrounded by fine sand, shown in Figure 1. The basalt blocks are of volcanic origin. The silicate sand occupies the space between the basalt blocks, and also fills in some of the surface vesicles, complicating the segmentation.

5.1.2 Type B: Massive basalt and lahar deposit

Outcrop exposure of massive basalt overlying older lahar deposit, shown in Figure 2. A lahar deposit, formed by violent pyroclastic flow, is visible in the lower portion of the image as a highly disordered accumulation of material with many irregular shapes and highly varying colour and texture. It is overlain by a layer of massive basalt, sourced from a later volcanic event. The basalt also has non-uniform colour, and some of the surface coatings in the basalt are of similar colour to the lahar. Fissures and ridges that intruduce linear shadows are also present.
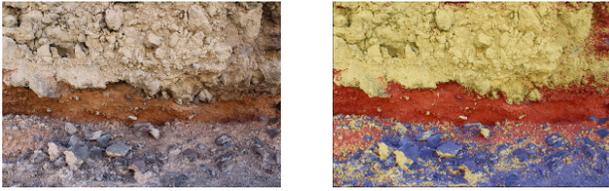
Figure 3. Example of scene type C. Layered volcanic materials, image number 0199, and its segmentation result.

5.1.3 Type C: Layered volcanic deposits

A succession of volcanic deposits from periodic events at the Cima volcanic flows, shown in Figure 3. Three layers of material are visible. Each has a different dominant colour, but each shows significant variation. Gradational changes are visible between the layers, and some blocks in the bottom layer have faces which are coated in material from the top layer. The boundaries are irregular, and enclaves of each material can be found within the others.

5.1.4 Type D: Complex intermixed volcanic materials

A complex scene, treated as three distinct classes, shown in Figure 4. Massive basalt overlies and is partially mixed into two other types of material, each of which is visibly heterogeneous. Such visibly complex scenes are common in a variety of geological settings.

5.1.5 Type E: Impact breccia

An exposed outcrop of breccia from the Sudbury impact structure, shown in Figure 7. Fractured by the violence of the impact, clasts of one type of rock are embedded in a matrix of another type. This type of material is common in impact craters and in volcanic settings.

5.1.6 Type F: Mineralized veins in Martian sandstone

An outcrop of clay-rich mudstone showing visible veins of calcium sulfate material [19], in the Yellowknife Bay locality of Gale Crater, Mars. This site was studied by the Mars Science Laboratory mission science team, and an example is shown in in Figure 10. It was selected as a high-priority science target, and a nearby outcrop of the same composition was the site of the first drill sample of the mission. The training image was acquired by the MSL rover's right Mastcam; the other two test images are of nearby exposures of the same material imaged by the left Mastcam, all on sol 133 of the mission. Significant and non-uniform coverage of reddish dust on both rock types, nearly ubiquitous on the Martian surface, complicates the vision problem.

Figure 4. Photograph of complex geological scene, image number 0495
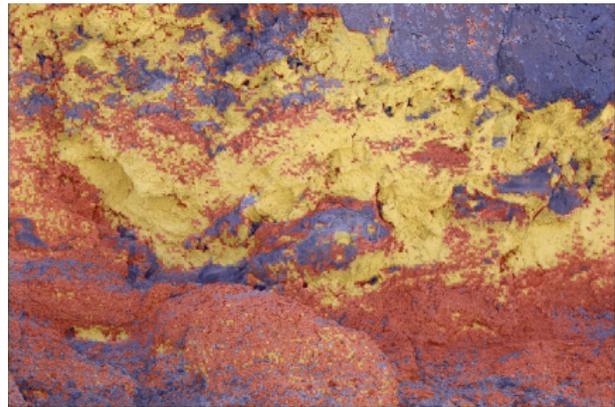
Figure 5. Segmentation map of image number 0495

5.2 Discussion

In general, the algorithm produces good quality results, both by reference to the Adjusted Rand Index figure of merit, and by visual inspection of segmentation maps. Even without the learned transformation matrix, the system produces results significantly better than a random pixel assignment, and in some types of scenes, far better. With the learning step included, the figure of merit increases in all cases, generally by a significant margin. The value of the training is also illustrated in Figure 6, which compares the separation of pixel vectors as plotted using a principal components analysis, and reprojected using the MDA-learned feature space representation. The projection is trained on a separate image. When applied to the new scene, it improves the alignment between k-means clusters and geologic unit classes. Black circles show the locations of cluster centroids.

The expanded feature space, including the MR8 filter bank, in most cases produces only a small improvement over the results with the basic feature set. The MR8 filter
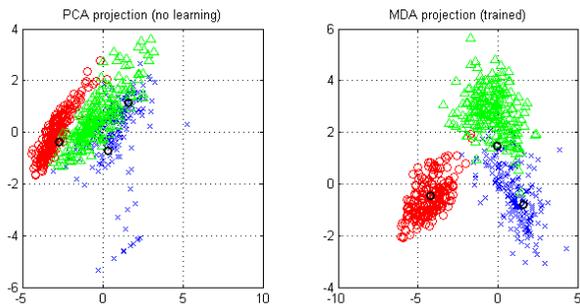
Figure 6. Normalized spread of the pixel vectors with the learned feature space representation from MDA, compared to an ordinary principal components analysis. Data from image #0497, trained on #0495, coloured by class label.

Figure 7. Photograph of Sudbury impact breccia, image number 9735

bank is designed to contribute information about visual texture, as is the gradient-derived channel in the basic feature space. Given the small marginal improvement of in the ARI in most cases, it appears that the gradient-derived channel adequately captures that textural information, for most purposes, and the high cost of computing the MR8 channels may not be justified in many applications.

As the scenes become more complex, with greater intra-class variation and complex boundaries between classes, the figures of merit are somewhat reduced. Nonetheless, even for very complex scenes, a visual inspection shows that the segmentation accurately reproduces the visual divisions that are salient to the human eye. An example of this is in scene type D, where three types of volcanic deposits are found in a single outcrop, with very complex mixing and irregular boundaries. Such a scene is shown in Figure 4, where a massive basalt, visible in the upper right hand corner, is intermixed with two different layers of volcanic material. The segmentation map for this image is shown in 5. The division of the scene by material type is evident, with the algorithm detecting even small isolated regions of one material embedded in another. Some stray pixels are apparent, attributed to the difficulty of training on this very difficult scene type. The Adjusted Rand Index for this segmentation is 0.742. While already likely adequate for a variety of follow-on uses such as steering a spectrometer instrument at the identified regions, futher improvements may be found by including new types of information in the feature space.

A particular demonstration of the value of the learning step is found in the case of the impact breccias from the Sudbury crater. Here the algorithm was, as usual, trained on a single example of a typical outcrop, then tested on other images of similar materials. One of these,
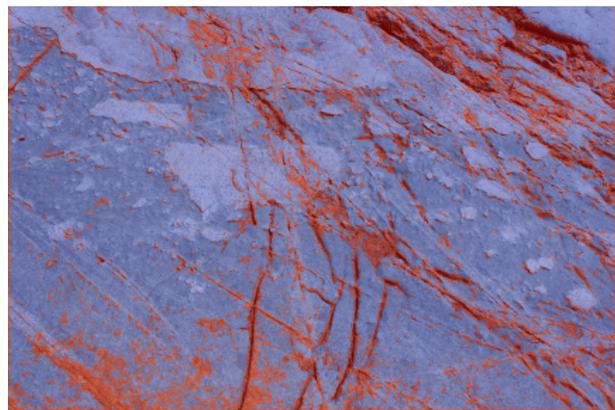
Figure 8. Segmentation map of image number 9735, without the learning step applied

Figure 9. Segmentation map of image number 9735, with the learning step applied
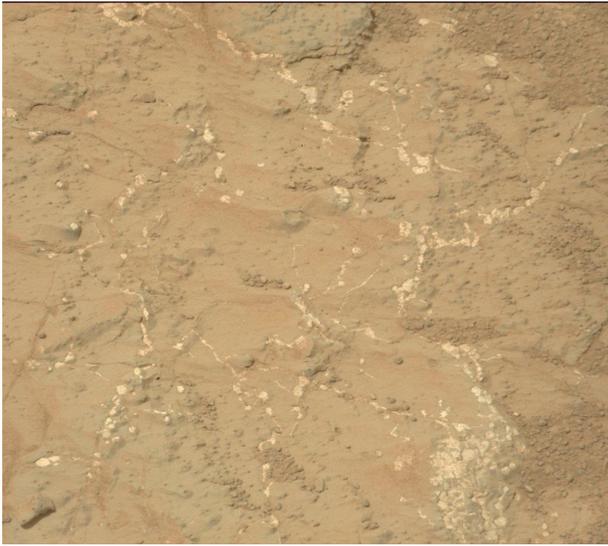
Figure 10. MSL MastCam-right photograph of calcium sulfate veins in mudstone, image reference s133r1. Image credit: NASA/JPL-Caltech/MSSS
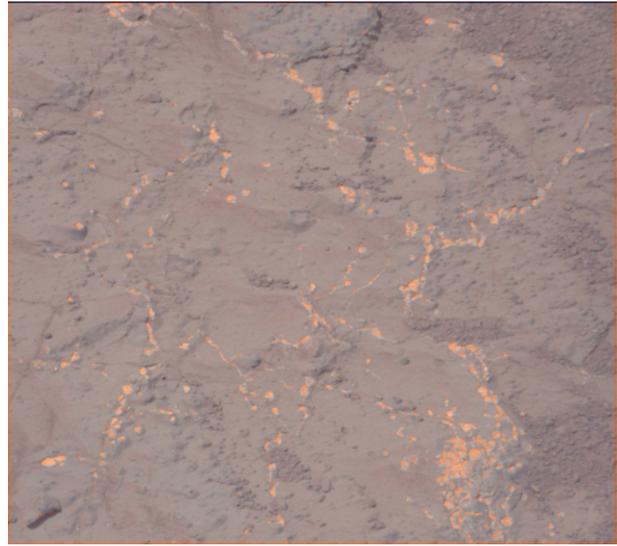
Figure 11. Segmentation map of image s133r1

with image number 9735, is shown in Figure 7. A reddish oxidation coating is visible on the surface of the outcrop. This uneven coating covers portions of both rock types (the darker-coloured ground mass, and the lighter-coloured clasts) in the breccia. This colouration is unrelated to the rock composition and potentially confusing to an algorithm relying on colour information, and without applying the learning step the results are unsatisfying. Figure 8 shows the segmentation map for the no-learning case, with ARI -0.072, slightly worse than random assignment. A naïve two-class clustering finds the numerically-significant difference between the smooth rock and dark shaded fractures. However, with the learned feature space transformation applied, the results are greatly improved, as shown in Figure 9, with the ARI rising to 0.793, even with the algorithm having been trained on a different image in which the oxidation coating was not present.

The learning step also improves the sedimentary example from Yellowknife Bay, Mars. Reddish-hued dust is nearly ubiquitous on rock outcrops on Mars, and like the oxidation coating in Sudbury, is potentially confusing to a vision algorithm. The segmentation is greatly improved by the training, even in the most challenging case (image s133ls3), where significant variation in dust cover makes the host rock appear more grey, rather than the dominant red hue seen in the training image in Figure 10. The challenging variation in colour from the dust cover is likely the reason why this is the geological setting in which the extended feature space, featuring the texture-rich MR8 filter bank, is most beneficial.

6 Conclusion

Further testing in new geologic settings is ongoing. In particular, the system is being tested as a means of detecting surface contaminations – recognizing dust partially covering a rock outcrop, for example, and discriminating it from the rock itself – for application to images of dusty contact spectrometry targets investigated by the MSL rover. Future developments of the technique could expand and optimize the feature space to account for optimal combinations and representations of the colour channels. An adaptation to multispectral imaging, as currently practiced with Mars-surface missions, is also planned. The system could also be adapted to other imaging modalities, depending on the instrumentation available in a given setting. As part of integration into a larger scheme for outcrop analysis, an autonomous method for determining the number of classes to use in the vector clustering is also in development.

Acknowledgments

References

[1] J. P. Grotzinger, J. Crisp, A. R. Vasavada, R. C. Anderson, C. J. Baker, R. Barry, D. F. Blake, P. Conrad, K. S. Edgett, B. Ferdowsi, R. Gellert, J. B. Gilbert, M. Golombek, J. Gómez-Elvira, D. M. Hassler, L. Jandura, M. Litvak, P. Mahaffy, J. Maki, M. Meyer, M. C. Malin, I. Mitrofanov, J. J. Simmonds, D. Vaniman, R. V. Welch, and R. C. Wiens, "Mars Science Laboratory mission and science investigation," *Space Science Reviews*, vol. 170, no. 1-4, pp. 5–56, September 2012. [Online]. Available: http://link.springer.com/article/10.1007/s11214-012-9892-2

[2] T. A. Estlin, B. J. Bornstein, D. M. Gaines, R. C. Anderson, D. R. Thompson, M. Burl, R. C. no, and M. Judd, "AEGIS automated science targeting for the MER Opportunity rover," *ACM Transactions on Intelligent Systems and Technology*, vol. 3, no. 3, pp. 50:1 – 50:19, May 2012. [Online]. Available: http://dl.acm.org/citation.cfm?id=2168764

[3] C. S. Cockell, G. R. Osinski, and P. Lee, "The impact crater as a habitat: Effects of impact processing of target materials," *Astrobiology*, vol. 3, no. 1, pp. 181–191, July 2004. [Online]. Available: http://online.liebertpub.com/doi/abs/10.1089/153110703321632507

[4] G. R. Osinski and E. Pierazzo, Eds., *Impact Cratering: Processes and Products*, 1st ed. Wiley-Blackwell, December 2012, ISBN: 978-1-4051-9829-5.

[5] V. C. Gulick, R. L. Morris, M. A. Ruzon, and T. L. Roush, "Autonomous image analyses during the 1999 Marsokhod rover field test," *Journal of Geophysical Research*, vol. 106, no. E4, pp. 7745 – 7763, April 2001. [Online]. Available: http://www.agu.org/pubs/crossref/2001/1999JE001182.shtml

[6] K. L. Wagstaff, D. R. Thompson, W. Abbey, A. Allwood, D. L. Bekker, N. A. Cabrol, T. Fuchs, and K. Ortega, "Smart, texture-sensitive instrument classification for in situ rock and layer analysis," *Geophysical Research Letters*, vol. 40, pp. 1–6, August 2013.

[7] V. Gor, R. Castaño, R. Manduchi, R. C. Anderson, and E. Mjolsness, "Autonomous rock detection for Mars terrain," in *Proceedings of AIAA Space 2001*, August 2001, albuquerque, New Mexico, 28 - 30 August 2001. [Online]. Available: http://users.soe.ucsc.edu/~manduchi/papers/AIAA01.pdf

[8] P. Viola and M. Jones, "Rapid object detection using a boosted cascade of simple features," in *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2001 (CVPR 2001)*, vol. 1, 2001, pp. I–511– I–518.

[9] D. R. Thompson and R. Castaño, "Performance comparison of rock detection algorithms for autonomous planetary geology," in *Proceedings of the 2007 IEEE Aerospace Conference*, 2007, big Sky, Montana, 3-10 March 2007.

[10] R. Castaño, R. C. Anderson, T. Estlin, D. DeCoste, F. Fisher, D. Gaines, D. Mazzoni, and M. Judd, "Rover traverse science for increased mission science return," in *Proceedings of the 2003 IEEE Aerospace Conference*, 2003, pp. 8–3629 – 8–3636, big Sky, Montana, 8-15March 2003. [Online]. Available: http://ieeexplore.ieee.org/xpl/articleDetails.jsp?tp=&arnumber=1235546

[11] T. Estlin, B. Bornstein, D. Gaines, D. R. Thompson, R. Castaño, R. C. Anderson, C. de Granville, M. Burl, M. Judd, and S. Chien, "AEGIS automated targeting for the MER opportunity rover," in *Proceedings of the 10th International Symposium on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS 2010), Sapporo, Japan, 29 August - 1 September 2010*, 2010. [Online]. Available: http://aegis.jpl.nasa.gov/publications/

[12] D. L. Bekker, D. R. Thompson, W. J. Abbey, N. A. Cabrol, R. Francis, K. S. Manatt, K. F. Ortega, and K. L. Wagstaff, "A field demonstration of a smart instrument performing autonomous classification of geologic surfaces," *Astrobiology*, 2014, in print.

[13] B. D. Bue, "An evaluation of low-rank Mahalanobis metric learning techniques for hyperspectral image classification," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing (JSTARS) (in press)*, 2013.

[14] J. V. Davis, B. Kulis, P. Jain, S. Sra, and I. S. Dhillon, "Information-theoretic metric learning," in *Proceedings of the 24th International Conference on Machine Learning*, 2007.

[15] N. Parrish and M. R. Gupta, "Dimensionality reduction by local discriminative gaussians," in *Proceedings of the 29th International Conference on Machine Learning*, 2012.

[16] J. Goldberger, S. Roweis, G. Hinton, and R. Salahutdinov, "Neighbourhood components analysis," in *Neural Information Processing Systems (NIPS) 2004*, 2004.

[17] M. Varma and A. Zisserman, "Texture classification: Are filter banks necessary?" *Proceedings of the 2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 2, pp. 691–698, June 2003.

[18] L. Hubert and P. Arabie, "Comparing partitions," *Journal of Classification*, vol. 2, pp. 193–218, 1985.

[19] D. T. Vaniman, D. L. Bish, D. W. Ming, T. F. Bristow, R. V. Morris, D. F. Blake, S. J. Chipera, S. M. Morrison, A. H. Treiman, E. B. Rampe, M. Rice, C. N. Achilles, J. P. Grotzinger, S. M. McLennan, J. Williams, J. F. Bell III, H. E. Newsom, R. T. Downs, S. Maurice, P. Sarrazin, A. S. Yen, J. M. Morookian, J. D. Farmer, K. Stack, R. E. Milliken, B. L. Ehlmann, D. Y. Sumner, G. Berger, J. A. Crisp, J. A. Hurowitz, R. Anderson, D. J. Des Marais, E. M. Stolper, K. S. Edgett, S. Gupta, N. Spanovich, and the MSL Science Team, "Mineralogy of a mudstone at yellowknife bay, gale crater, mars," *Science*, vol. 343, no. 6169, p. 8, January 2014. [Online]. Available: http://www.sciencemag.org/content/343/6169/1243480.full