Mining Remote Image Repositories with Application to Mars Rover Stereoscopic Image Datasets

Andrew Willis, Waseem Shadid, Martha Cary Eppes
University of North Carolina at Charlotte
9201 University City Blvd., Charlotte, NC 28223

ABSTRACT

As of December 2008, the two Mars rover spacecraft Spirit and Opportunity have collected more than 4 years worth of data from nine imaging instruments producing greater than 200k images which includes both raw image data from spacecraft instruments and images generated by post-processing algorithms developed by NASA's Multimission Image Processing Laboratory (MIPL). This paper describes a prototype software system that allows scientists to browse and data-mine the images produced from NASA's Mars Exploratory Rover (MER) missions with emphasis on the automatic detection of images containing rocks that are of interest for geological research. We highlight two aspects of our prototype system: (1) software design for mining remote data repositories, (2) a computationally efficient image search engine for detecting MER images that containing rocks. Datatype abstractions made at the software design level allow users to access and visualize the source data through a single simple-to-use interface when the underlying data may originate from a local or remote image repository. Data mining queries into the MER image data are specified over chronological intervals denoted (sols) as each interval is a solar day. As in other mining applications, an automatic detection and classification algorithm is used to compute a relevance score that represents how relevant a given recorded image is to the user-specified query. Query results are presented as list of records, sorted by their relevance score, which the user may then visualize and investigate to extract information of interest. Several standard image analysis tools are provided for investigation of 2D images (e.g., histogram equalization, edge detection, etc.) and, when available, stereoscopic data is integrated with the image data using multiple windows which show both the 2D image and 3D surface geometry. The combination of data mining and a high-quality visualization interface provides MER researchers unprecedented access to the recorded data.

1. INTRODUCTION

Images from the Mars Rovers *Spirit* and *Opportunity* have been streaming to the Earth since their landing on the Martian surface in January 2004. Datasets are periodically released to the public as a collection of image files organized by the number of Earth days (*sols*) since the landing where each day commonly includes several hundred images. Current estimates indicate that there are roughly 200k images available from the Mars Exploratory Rover (MER) missions. These images and their associated 3D data represent an enormous bank of information. In particular, these data have the potential to provide three-dimensional geological field data (strikes, dips, dimensions) which are basic, yet imperative to addressing a wide range of geological hypothesis related to fluvial dynamics, impacts, dune migration, tectonics and more. Earth Science workers studying Mars commonly use large-sample empirical relationships to confirm or refute scientific hypotheses (e.g. [1]). However, currently the average scientist cannot readily access or derive statistically significant scientific inferences from MER data. With few exceptions (e.g. [2–4]), no significant dimensional or orientation data has been presented from MER missions to-date, despite the wealth of imagery. This lack of data analysis is likely due to the fact that no software exists to readily access MER data without painstakingly downloading the files and running platform specific, special purpose software such as NASA's ISIS and NASAVIEW programs. This article presents a program capable of (i) browsing such repositories online, (ii) visualizing the uncompressed 2D image data and, for stereo image pairs, the 3D image data online and (iii) dispatching data mining tasks on remote repositories that locate rocks within the MER images.

Further author information: (Send correspondence to A. Willis)
A. Willis: E-mail: arwillis@uncc.edu, Telephone: 1 704 687 8420

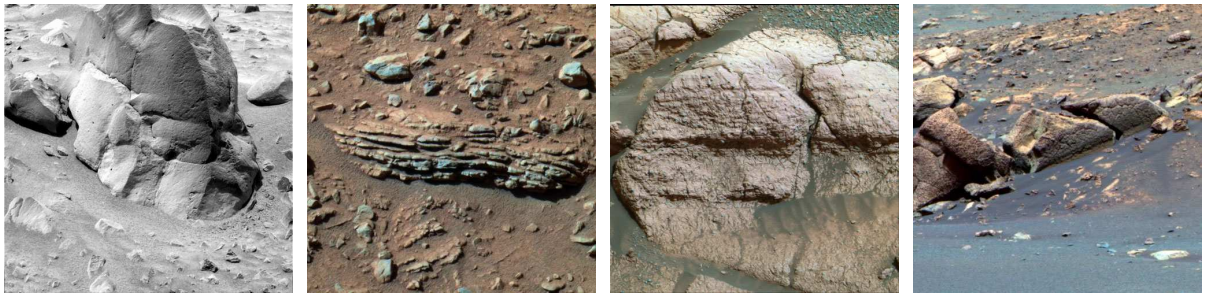

(a) (b) (c) (d)

Figure 1. Exemplar Pancam Images: (a) Spirit: A rock dubbed "Humphrey." (released Feb. 27, 2004) (b) Spirit: This is a false-color composite image of the horizontally layered rock dubbed "Tetl" at the "Columbia Hills" in Gusev Crater. (released October 7, 2004) (d) Opportunity: Scientists believe these sphere-like grains at the rock outcrop were formed in pre-existing wet sediments. (released March 18, 2004) (e) Opportunity: A false-color image of a rock informally named "Earhart" on the lower slopes of "Endurance Crater." The rock was named after the pilot Amelia Earhart. Like other rocks dotting the bottom of Endurance, scientists believe fractures in Earhart could have been formed by impact and water processes. (released October 7, 2004) Pictures and descriptions courtesy of D. Savransky and J. Bell (Cornell) / JPL / NASA [5].

2. RELATED WORK

NASA currently has a large suite of software developed in-house for the analysis of all mission data. Through several discipline-related 'nodes' NASA offers customized software packages for PDS (Planetary Data System) record viewing and processing. This software includes BulkDownloader, NASAVIEW, and ISIS (Integrated System for Imagers and Spectrometers) for general data processing. Some tools such as MER Analyst Notebook and Planetary Image Analysis are specifically designed to aid the user in visualizing MER images. All of these packages are available, although not uniformly, for varying OS platforms or through the web. Despite the availability of this software some drawbacks still exist for users that wish to extract more than the visual information provided through JPEG versions of MER images. MER Analyst Notebook (MERAN), for example, is the most obvious portal to MER data for an average non-NASA user. MERAN allows the user to browse through MER data by visualizing all images by *sol*. Once an image is chosen, data accompanying that image are provided as links which can then be downloaded. As such, each image and its accompanying data must be downloaded individually. MER Analyst Notebook does not provide 3D image analysis capabilities, however, so then the user must download, install, and access additional software separately in order to perform research. Many of these systems, in turn, require platform specific software applications and that large databases be maintained by the user. These tasks alone represent significant hurdles for the non-computer savvy scientist, teacher or general user to overcome, and therefore, in many cases, likely preclude any complex or significant exploitation of the data. The proposed system shares in-part functionality from many different PDS-offered applications, however, the major impetus for the proposed software is to facilitate automated processing and analysis of PDS data records. In doing so, we hope to greatly increase the accessibility and functionality of the NASA data for the scientific, teaching, and hobby-scientist community.

In [6], Frawley et. al. describe data mining as "the nontrivial extraction of implicit, previously unknown, and potentially useful information from data." The abundance of imaging systems available today are generating unprecedented amounts of data. In many cases researchers need specific information from a small subset of the recorded imagery. Rather than painstakingly reviewing each image manually, data mining algorithms automatically process images to determine those relevant to the researchers need. Implementation of such systems require image processing and pattern recognition algorithms which compute features and classify pixels within images to generate relevance scores that is the system's characterization of how relevant an given image is to a researchers query. Many papers have investigated methods for extracting image features such as wavelets [7], vocabulary trees [8], hierarchical K-means dictionaries [9], and bag-of-features methods [10]. For geological features, there are several methods that have been proposed in the literature. In [11], the authors present a 6-step method for detection of geological features within images from a robotic rover that uses multiple input images of the same surface region and a Bayesian Belief Network to classify rocks within images. Other related works on computing geological features in from MER rover images include . The methods proposed therein often make use of color or stereoscopic data which is time consuming to generate (see Figure 1(b-d) for examples of such images and [12, 13] for details on how these images are generated) if these images are not already available and, in cases where the data is available, the

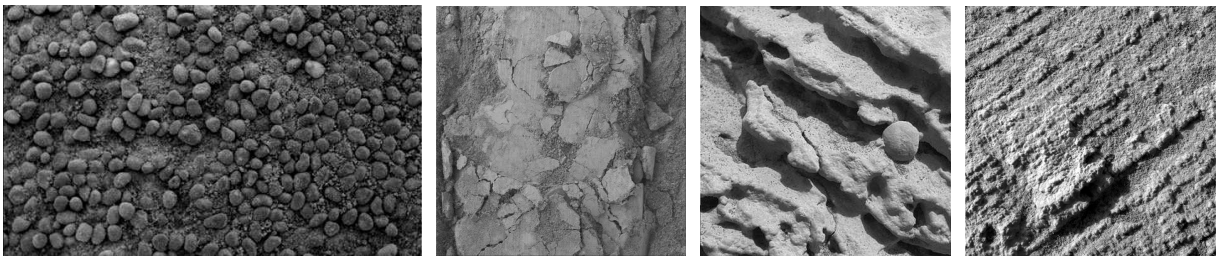

(a) (b) (c) (d)

Figure 2. Exemplar MI Images: (a) Spirit: A close-up view of the coarse grains on the crests of the waves of windblown soil in Gusev Crater. (released Feb. 26, 2004) (b) Spirit: This close-up view of a patch of soil on Mars was taken by Spirit's Microscopic Imager after the soil was compacted by one of the rover's wheels. The area shown is approximately 3 centimeters (1.2 inches) square. (released November 29, 2004) (c) Opportunity: A close-up view of fine grained rock seen in the rock outcrop at Meridiani Planum by Opportunity's Microscopic Imager. (released Feb. 9, 2004) (d) Opportunity: A microscopic image of "Flatrock," taken on Opportunity's 43rd sol at Meridiani Planum. The very fine layered nature of "Flatrock" preserves a record of sediment accumulation that occurred long ago. Layers are only a few millimeters thick and are accentuated by a distinctive mineral composition that makes them more resistant to erosion. (released March 8, 2004). Pictures and descriptions courtesy of NASA/JPL/Cornell/USGS [5].

additional information can significantly increase the processing time required to generate results for general queries over a very large number of images.

3. MER IMAGING SYSTEMS AND MISSION DATA RECORDS (EDRS & RDRS)

There are nine 2D-imaging systems on *Spirit* and *Opportunity* which are jointly responsible for generating approximately 50-350 recorded images per day, per rover. These images come from two systems on each of the robots and are listed below:*

- Science instruments, aka Athena Package [14]
  - Panoramic cameras (2) (see Figure 1)
  - Microscopic imager (MI) (see Figure 2)
- Cameras for Hazard Avoidance and Navigation
  - Dual front hazcams (hazard detection and avoidance cameras) (2)
  - Dual rear hazcams (2)
  - Dual navcams (navigational cameras) (2)

As of Dec. 18th, 2008 there are 1530 sols available to the public through NASA's Planetary Data System (PDS) [15] which constitutes the raw source data for our data mining system.

Each recorded image is transferred as part of the rover telemetry to Earth and reconstructed as a raw image referred to as an Experimental Data Record (EDR), e.g., Figures 1(a) and 2(a-d). EDR images are often processed using a suite of software tools developed by the Multimission Image Processing Laboratory (MIPL) a part of the NASA Jet Propulsion Laboratory (JPL). These programs take as input one or more EDR images and integrate the measured data into a new output image referred to as a Reduced Data Record (RDR). Each EDR image may be processed to generate up to 12 different types of RDR images but typically between 4-8 RDRs may be generated for a new EDR. Reduced data records are numerous and include color images formed from multiple images of a surface using different optical filters [12, 13], e.g., Figure 1(b-d), 3D stereoscopic images reconstructed from the Pancam stereo system, e.g., Figure 5(b), and image mosaics, e.g., Figure 3.

*We list only those instruments that generate 2D images below. The Athena Package also includes 5 additional non-image sensors.

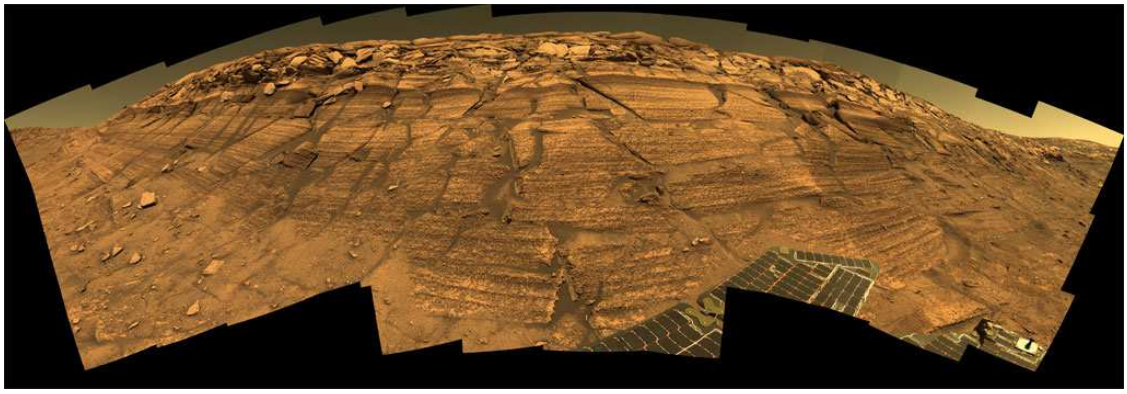

Figure 3. An image mosaic from Opportunity: This approximate true-color mosaic of "Burns Cliff" in "Endurance Crater" is a composite of 46 different images, each acquired in seven different Pancam filters. (released December 13, 2004). Picture and description courtesy of NASA/JPL/Cornell/USGS [5].

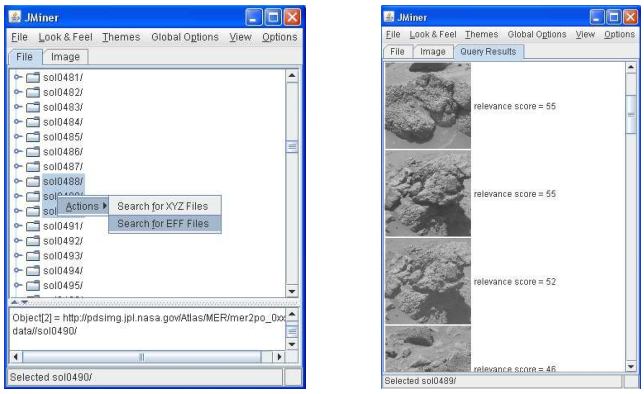

Figure 4. (a) The data mining interface shows the list of available sols for the selected rover. The user may highlight any collection of sols to issue the data mining query. In the example above, the user has selected three sols (488,489,490). (b) Shows the tabulated results for each query which are updated as the search proceeds. Images are listed in order of decreasing relevance.

Of particular interest for geological research is the 3D stereoscopic data produced by the Pancam imaging system. This data is generated via a technique known as stereoscopic 3D reconstruction and is stored in a XYZ Reduced Data Record (RDR) which contains estimates for 3D (x,y,z) surface positions of objects viewed within a stereoscopic pair of Pancam EDR images [16]. The underlying concept of 3D reconstruction from images taken from a pair of cameras is that by knowing or estimating the image formation properties of each camera, their relative pose, and the pixel pair in each digital image that corresponds to a specific 3D surface location, one may invert the image formation process and find the 3D locations responsible for reflecting the light sensed by the camera.

4. DATA MINING INTERFACE

Our data mining system is implemented in Java and can be launched from any Java-enabled computer. Available remote files are visualized by the user and may be selected interactively to initiate searches over large file groups (see Figure 4). The interface shows the MER PDS data available on the remote server which is organized as a hierarchical tree structure where each solar day since the mission start date is listed numerically, e.g., sol0010 corresponds to solar day 10. The user selects those aspects of the PDS tree structure that are of interest for searching as shown in Figure 4(a). Once the appropriate sources of data have been indicated, the user can initiate a query which, for our initial system, seeks rocks within EDR images. In a separate thread of execution, our rock detection algorithm is executed which computes relevance scores for EDR images within the selected data sources and posts the results to a query window within the program interface (see Figure 4(b)).

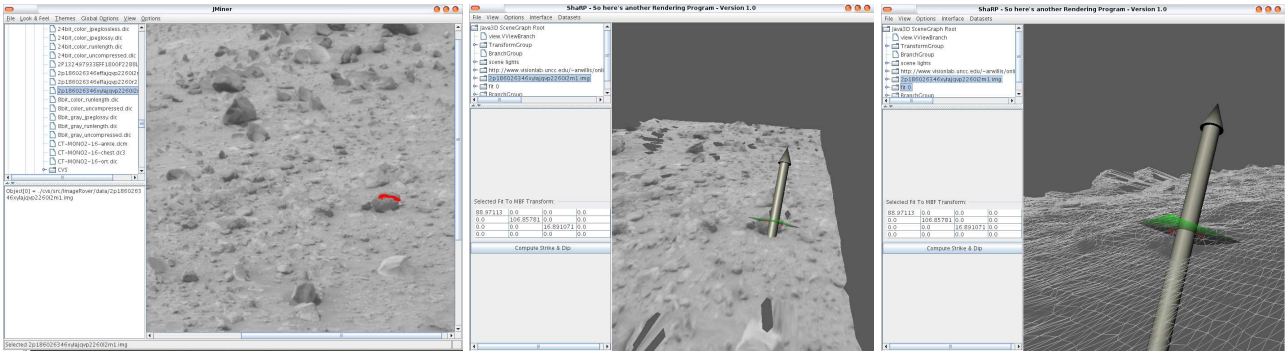

(a) (b) (c)

Figure 5. (a,b,c) are all screen captures of the described program. (a) shows a single 2D image from a stereoscopic pair taken of the Martian surface (2p186026346xylajqvp2260l2m1.img). A curve (in red) has been manually drawn into the image by the user to denote the path of a crack on a Martian rock of interest. (b,c) show two views of the corresponding 3D space curve in the stereoscopic image data and a cylinder indicating the normal of the best-fit plane to the 3D space curve, i.e., the approximate space-curve orientation.

5. VISUALIZATION AND ANALYSIS OF DATA RECORDS

Search results may be reviewed visually and, for images containing features of specific interest, a suite of analysis tools is provided for both 2D EDR images and, when available, 3D stereoscopic RDR images. The provided tools focus on enabling users to quickly make accurate measurements of the most common measurements made in the field by geologists which are the size, orientation and the distance between objects or geological features. These measurements are universal to a broad variety of geological applications and are critical to the successful analysis and study of Mars' surface and its associated geological processes. For this reason, our prototype system allows search results to be examined in both 2D and 3D visualizations simultaneously. This interface is particularly useful for rock geometry analysis. For example, one may extract the geometry of a rock fracture by tracing the image of the fracture path in a 2D stereoscopic image with a mouse (see Figure 5a). The software then back-projects this 2D curve into the 3D stereoscopic data to provide detailed 3D geometric information regarding the geometry (length, width and orientation, etc.) of the fracture which are of significance for geological research (Fig. 5b).

6. MARTIAN GEOLOGY: WHAT NEW QUESTIONS CAN THIS SYSTEM HELP ANSWER?

There are countless geological hypotheses that might be tested for the Mars surface with the type of simple geometric data that can be readily acquired with our system. For example the thickness and orientation of bedding planes such as those visible in (c) and (d) of both Figures 1 and 2 might provide insight into past aeolian or fluvial processes acting on the Mars surface. Measuring the orientation of these bedrock features might also allow geo-scientists to test hypotheses related to possible past plate tectonics on Mars (e.g. [17]). The size and distances between stones on surfaces such as that visible in Figure 1(b) might provide data necessary to analyze ejecta emplacement. Yet, the impetus for the development of this prototype software was pursuit of single geological test-problem related to the physical weathering of surface rocks on Mars.

Mechanical weathering is a process which affects rocks on Mars. Rocks visible in MER images are commonly characterized by sub-vertical fractures, and small rocks often exhibit identical spectral properties to adjacent larger rocks suggesting that they have been mechanically separated [18]. The primary processes responsible for the fracturing of rocks on Mars is not apparent, however. Despite the availability of high-resolution images of rocks present on Mars's surface, there has been little if any investigation into the physical processes by which these rocks deteriorate. Hypotheses include breakdown associated with original ejecta emplacement [18] and salt weathering by thermal expansion or hydration (e.g. [19]). Recent work described in [20], proposes that Martian fractures originate from tensile stresses that arise as rocks are directionally heated and cooled during the sun's east-to-west transit across the sky. As crack orientations should reflect the orientation of the largest of these recurrent thermal stresses, this hypothesis can be tested for Martian rocks by measuring the orientations of cracks from MER mission products. In the absence of strong mechanical anisotropy, the majority of cracks should be oriented north-south, as were the majority of the >700 cracks measured in a study of rocks on Earth's deserts [20].

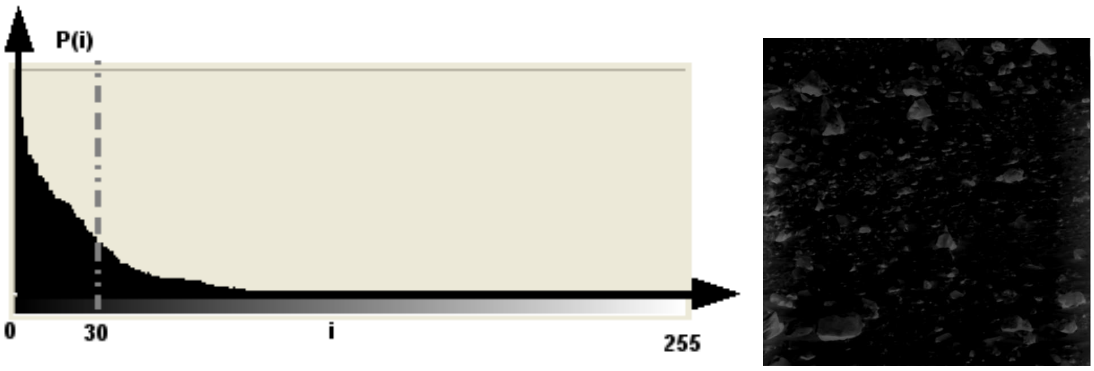

(a) (b)

Figure 6. Our simplistic algorithm inspects the intensity variations of the EDR image from a local estimate of the image background. (a) Shows the distributions of intensities for a collection of Martian images after background subtraction, i.e., the histogram of $\Delta I_{bg}(x, y)$ from equation (3) for several exemplar images. (b) Shows an example of an EDR image after subtracting this local estimate of the background, i.e., $\Delta I_{bg}(x, y)$.

If other processes such as salt hydration are responsible for mechanical weathering, then the fractures will have random orientations.

Critical elements required to test the proposed hypotheses include: 1) non-biased selection of images and rocks with cracks and 2) a statistically substantial data set of crack orientations (>1000) accompanied with rock dimension information. Our work began by using the current NASA-sponsored web-based interface to find digital camera images which included rock fractures and collars. This involved a manual search for cracked rocks by individually examining the many thousands of images available. This process became prohibitively expensive in terms of time and resources, and it became readily clear that this method was not going to produce statistically significant results given the allocated time and economic resources. This situation is likely not unique. Most Earth Science applications of MER data require similar examination (e.g. [21]). Earth scientists studying Mars commonly use large-sample empirical relationships to confirm or refute scientific hypotheses (e.g. [1]).

Critical elements required to test the proposed hypotheses include: 1) non-biased selection of images and rocks with cracks and 2) a statistically substantial data set of crack orientations (>1000) accompanied with rock dimension information. Both of these elements can be provided by an effective data mining system which is the impetus for this work.

7. ROCK DETECTION ALGORITHM

The algorithm for rock detection is discriminative. It is intended to work quickly and exploit knowledge we have obtained by examining a number of images taken by the rovers of the Martian surface. Specifically, we begin our explanation of the algorithm by examining the average intensity of the Martian soil and of rocks within MER EDR images. We found that averaging over large image regions provided an estimate of the Martian soil intensity sufficient for coarse identification of rocks within EDR images. We then define an algorithm based on the expected variation of rock highlights and shadows away from an estimate of the Martian soil intensity computed by averaging over large regions of the image data. The histogram of the absolute value of these variations is shown in Figure 6(a) for a selected set of exemplar images from the MER image repository.

The simplicity of this approach allows for efficient computation of rock detection within images which is important for mining very large image repositories such the the MER image repository. The algorithm steps are:

1. Use median filtering to remove image noise, a standard image processing algorithm as described in [22].

2. Approximate the intensity of the background at each pixel location by computing the average pixel intensity over a large $M \times N$ region around each pixel as in equation (1).

$$I_{bg}(x,y) = \frac{1}{NM} \sum_{n=\frac{-N+1}{2}}^{\frac{N+1}{2}} \sum_{m=\frac{-M+1}{2}}^{\frac{M+1}{2}} I(x,y), \quad N, M > 0 \, and \, odd \tag{1}$$

A global image feature, σ, is computed from $I_{bg}(x,y)$ which depends upon the expected intensity variations between Martian soil and Martian rocks as represented by the empirical distribution of each intensity variation, $P(i)$, that is computed by taking a set of exemplar MER images containing rocks and computing the histogram of observed intensity values in the images$\Delta I_{bg}(x,y)$ over the exemplar set (see equation (3)). Our empirical distribution , $P(i)$, is shown in Figure 6(a). The feature σ is a modified version of standard deviation computed as shown in equation (2) that weights each deviation by the probability of observing the intensity variation, $P(i)$.

$$\sigma = \sum_{i=0}^{255} \left(P(i) \left| I(x,y) - I_{bg}(x,y) \right|^2 \right)^{\frac{1}{2}} \tag{2}$$

3. Perform background subtraction as shown in equation (3), see Figure 6(b) for an example of an image $\Delta I_{bg}(x,y)$.

$$\Delta I_{bg}(x,y) = |I(x,y) - I_{bg}(x,y)| \tag{3}$$

4. Classify each pixel by assigning each pixel to a class Ω such that each pixel is either a member of the background, $\Omega = 0$, or a candidate rock pixel, $\Omega = 1$. Classification is accomplished by thresholding $\Delta I_{bg}(x,y)$ with a threshold value $\lambda = 5$ for our results as in equation (4)

$$\Omega(x,y) = \begin{cases} 0 & if \quad \Delta I_{bg}(x,y) > \lambda \\ 1 & if \quad \Delta I_{bg}(x,y) \leq \lambda \end{cases} \tag{4}$$

5. From the binary image $\Omega(x,y)$, a set of connected closed regions with labels $\Omega = 1$ are extracted. We denote the i^{th} closed image region $\mathbb{R}_i$. We then compute area of each region in pixels denoted, $N_{\mathbb{R}_i}$, and the length of the boundary of each region (in pixels) denoted, $C_{\mathbb{R}_i}$. Those regions whose size or perimeter is smaller than the user-specified minimum values are discarded. The minimum rock area, and boundary length for our results are 15 pixels and 225 pixels respectively.

6. We then compute the feature $\sigma_{\mathbb{R}_i}$ as shown in equation (5). This value determines the salience of each rock region in the image based the variation of the rock intensities from the background image I_{bg}. Regions whose $\sigma_{\mathbb{R}_i}$ value is less than a user-specified minimum value are discarded. Our results are generated using a minimum $\sigma_{\mathbb{R}_i}$ value of 30.

$$\sigma_{\mathbb{R}_i} = \frac{1}{N_{\mathbb{R}_i}} \sum_{(x,y) \in \mathbb{R}_i} \Delta I_{bg}(x,y) \tag{5}$$

7. Let N_{rocks} denote the number of rock regions detected in step (5). The image relevance score, R, is computed as a value $R \in [0, 100]$ that depends on the number of rocks and the average salience of the rocks within the image as indicated in equation (6) where $E[N_{rocks}]$ is the expected number of rocks in an image of interest and is specified by the user. Our results are generated using $E[N_{rocks}] = 5$.

$$R = 0.2 \max\left(100, 100\frac{N_{rocks}}{E[N_{rocks}]}\right) + 0.8\frac{1}{N_{rocks}} \sum_{i=1}^{N_{rocks}} \max\left(100, \frac{100\sigma_{\mathbb{R}_i}}{\sigma}\right) \tag{6}$$

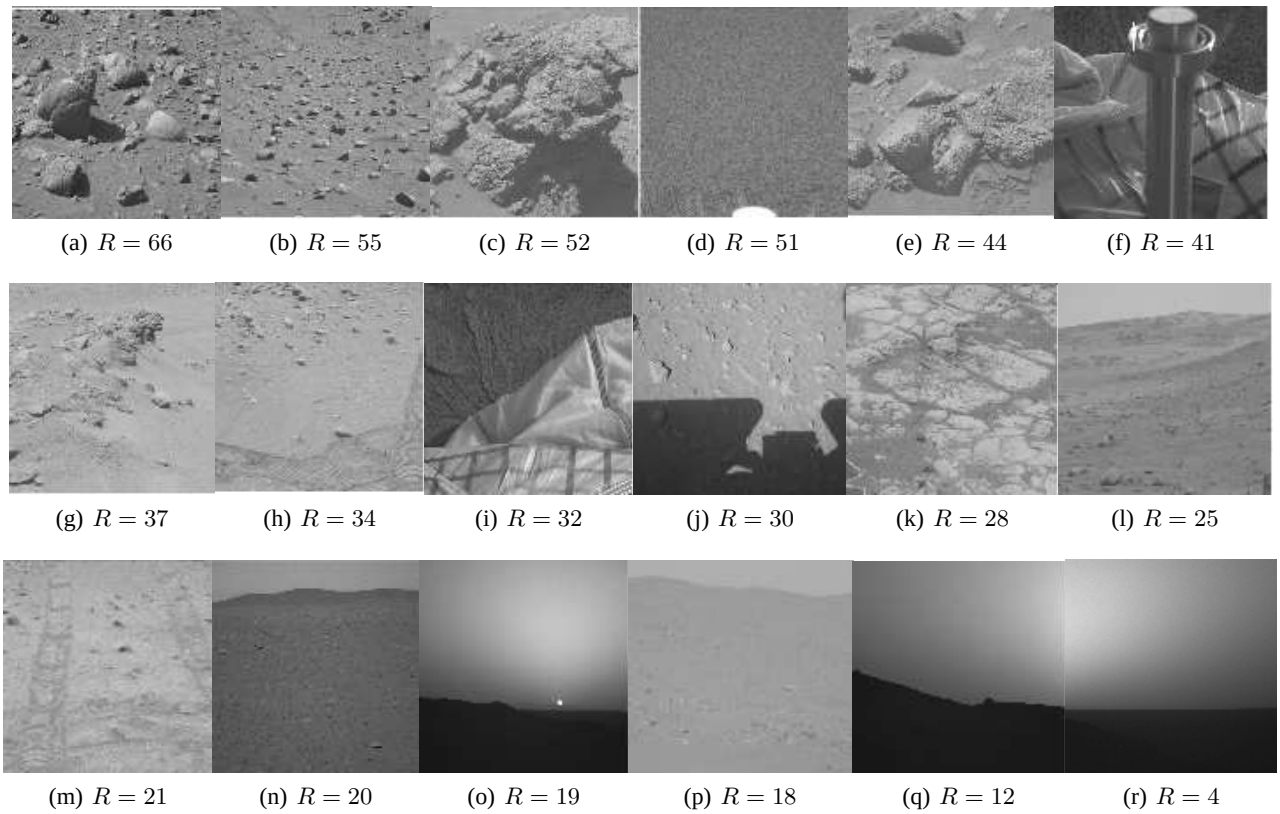

(a) $R = 66$ (b) $R = 55$ (c) $R = 52$ (d) $R = 51$ (e) $R = 44$ (f) $R = 41$

(g) $R = 37$ (h) $R = 34$ (i) $R = 32$ (j) $R = 30$ (k) $R = 28$ (l) $R = 25$

(m) $R = 21$ (n) $R = 20$ (o) $R = 19$ (p) $R = 18$ (q) $R = 12$ (r) $R = 4$

Figure 7. (a-r) show EDRs from data mining searches sorted by their relevance value, R. Note that we have purposefully inserted three results which we consider to be erroneous: (d),(f), and (i). Which are images that either contain trivial intensity information (d) or image structures other than rocks (f) and (i). Preliminary empirical results show a good agreement between the order of the returned images in terms of their relevance score and those images containing rocks to be found by the query. We are currently working towards validating these results with geologists.

8. RESULTS

Considering the computational simplicity of our algorithm, we are pleased with the results it provides. In cases where clear and distinguishable rocks images are available, the algorithm produces high relevance's scores. In cases where the observed rock intensities are very close to the background intensities low relevance score values are produced. Figure 7 shows a nearly uniform distribution of the observed results. We have purposefully inserted more false-positives than is typically encountered to provide the reader some insight on cases when the algorithm provides unexpectedly high relevance scores.

For the settings specified in §7, we see that images that have rocks generally have a relevance score above 30. Images that contain rocks that have intensity values close to background values or have areas large enough to dominate the background average filter generally score between 20 and 30. Images that do not have rocks or have rocks that do not match our settings generally score less than 20.

9. CONCLUSION

This article describes a system for performing data mining operations on remote image repositories. The implemented system is used to locate MER images that contain rocks of interest to geologists. The data mining interface allows users to define searches over an arbitrary subset of available MER data and to control the salience, area, and boundary length of rocks to be detected within the MER images. Query results may be analyzed using a 2D image interface and, when stereoscopic XYZ RDR data is available, an integrated 3D interface. A major benefit of remote data mining is that it allows scientists to discover and analyze data of relevance to their scientific investigations without needing the resources to store and manage the data. Query results are provided as a list of images sorted by a relevance score computed by an image

processing algorithm that automatically estimates how relevant each image is to the user-specified query. The algorithm is both simple, computationally efficient and provides promising results which allows for large queries to be processed in an acceptable amount of time. Planned extensions to this system seek to bolster the query and analysis interface and improve upon the performance of the system in terms of network bandwidth, query result computation time, and accuracy of computed relevance values. In conjunction with modifications to our relevance computations, we also seek to validate the computed relevance values with the help of practicing geologists.

REFERENCES

1. J. A. Grant, R. Arvidson, J. F. Bell III, N. A. Cabrol, M. H. Carr, P. Christensen, L. Crumpler, D. J. Des Marias, B. L. Ehlmann, J. Farmer, M. Golombek, F. D. Grant, R. Greeley, K. Herkenhoff, R. Li, H. Y. ScSween, D. W. Ming, J. Moersch, J. W. Rice Jr., S. Ruff, L. Richter, S. Squyres, R. Sullivan, and C. Weitz, "Surficial deposits at gusev crater along spirit rover traverses.," *Science*, vol. 305, pp. 807–810, August 2004.
2. K. W. Lewis, O. Aharonson, J. P. Grotzinger, S. W. Squyres, J. F. Bell III, L. S. Crumpler, and M. E. Schmidt, "Structure and stratigraphy of home plate from the spirit mars exploration rover," *Journal Geophysical Research*, vol. 113, 2008. E12S36, doi:10.1029/2007JE003025.
3. M. P. Golombek, R. E. Arvidson, J. F. Bell III, J. A. Christensen, P. R. Crisp, L. S. Crumpler, B. L. Ehlmann, R. L. Fergason, J. A. Grant, R. Greeley, A. F. C. Haldemann, D. M. Kass, T. J. Parker, J. T. Schofield, S. W. Squyres, and R. W. Zurek, "Assessment of mars exploration rover landing site predictions," *Nature*, no. 436, pp. 44–48, 2005.
4. M. Golombek, L. S. Crumpler, J. A. Grant, R. Greeley, N. A. Cabrol, T. J. Parker, J. W. Rice Jr., J. G. Ward, R. E. Arvidson, J. E. Moersch, R. L. Fergason, P. R. Christensen, A. Castaño, R. Castaño, A. F. C. Haldemann, R. Li, J. F. Bell III, and S. W. Squyres, "Geology of the gusev cratered plains from the spirit rover transverse," *Journal of Geophysical Research*, vol. 111, 2006.
5. D. Savransky and J. B. C. . J. . NASA, "Nasa jet propulsion laboratory mars rover program," 2004.
6. *Knowledge Discovery In Databases*. AAAI Press/MIT Press, Cambridge, MA., 1996.
7. T. Li, Q. Li, S. Zhu, and M. Ogihara, "A survey on wavelet applications in data mining," *SIGKDD Explor. Newsl.*, vol. 4, no. 2, pp. 49–68, 2002.
8. D. Nistér and D. Stewénius, "Scalable recognition with a vocabulary tree," in *Conf. on Computer Vision and Pattern Recognition (CVPR)*, vol. 2, pp. 2161–2168, 2006.
9. S. Lazebnik, C. Schmid, and J. Ponce, "A sparse texture representation using local affine regions," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 27, no. 8, pp. 1265–1278.
10. E. Nowak, F. Jurie, and B. Triggs, "Sampling strategies for bag-of-features image classification," in *European Conference on Computer Vision (ECCV)*, vol. IV, pp. 490–503, 2006.
11. D. Thompson, T. Smith, and D. Wettergreen, "Data Mining During Rover Traverse: From Images to Geologic Signatures," in *'i-SAIRAS 2005' - The 8th International Symposium on Artificial Intelligence, Robotics and Automation in Space*, vol. 603 of *ESA Special Publication*, August 2005.
12. J. F. Bell III, S. W. Squyres, K. E. Herkenhoff, J. N. Maki, H. M. Arneson, D. Brown, S. A. Collins, A. Dingizian, S. T. Elliot, E. C. Hagerott, A. G. Hayes, M. J. Johnson, J. R. Johnson, J. Joseph, K. Kinch, M. T. Lemmon, R. V. Morris, L. Scherr, M. Schwochert, M. K. Shepard, G. H. Smith, J. N. Sohl-Dickstein, R. Sullivan, W. T. Sullivan, and W. M., "The mars exploration rover athena panoramic camera (pancam) investigation," *J. Geophys. Res.*, vol. 108, 2003. E12, doi:10.1029/2003JE002070.
13. J. F. Bell III, J. R. Joseph, J. Sohl-Dickstein, H. Arneson, M. Johnson, M. Lemmon, and D. Savransky, "In-flight calibration of the mars exploration rover panoramic camera instrument," *J. Geophys. Res.*, vol. 111, 2006. E02S03, doi:10.1029/2005JE002444.
14. "Mars exploratory rover athena scientific instrument package," 2004. http://athena.cornell.edu/.
15. D. Savransky and J. Bell III, "Nasa planetary data system - jet propulsion laboratory (cornell) / jpl / nasa."
16. L. Matthies, M. W. Maimone, A. E. Johnson, Y. Cheng, R. G. Willson, C. Villalpando, S. B. Goldberg, A. Huertas, A. Stein, and A. Angelova, "Computer vision on mars," *International Journal of Computer Vision (IJCV)*, vol. 75, no. 1, pp. 67–92, 2007.
17. N. H. Sleep, "Martian plate tectonics," *Journal Geophysical Research*, vol. 99, 1994.
18. G. Greely, S. Squyres, *et al.*, "Wind-related processes detached by the spirit rover at gusev crater, mars.," *Science*, vol. 305, pp. 810–821, August 2004.

19. G. Wells and J. Zimbelman, “Extraterrrestrial arid surface processes,” *Arid Zone Geomorphology*, pp. 335–358, August 1989.
20. L. McFadden, M. Eppes, and B. Hallet, “Physical weathering in arid landscapes due to diurnal variation in solar heating,” *Geological Society of America Bulletin*, vol. 117, pp. 161–173, 2005.
21. S. Squyres, R. Arvidson, *et al.*, “The spirit rover’s athena science investigation at gusev crater, mars.,” *Science*, vol. 305, pp. 794–799, August 2004.
22. E. Trucco and A. Verri, *Introductory Techniques for 3D Computer Vision*. Prentice Hall, 1998.