

Performance Comparison of Rock Detection Algorithms for Autonomous Planetary Geology

David R. Thompson¹ and Rebecca Castaño²

¹The Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15206
(412) 628-7413
drt@ri.cmu.edu

² Jet Propulsion Laboratory
California Institute of Technology
Pasadena, CA 91109
(818) 393-3544
Rebecca.Castaño@jpl.nasa.gov

Abstract—Detecting rocks in images is a valuable capability for autonomous planetary science. Rock detection facilitates selective data collection and return. It also assists with image analysis on Earth. This work reviews seven rock detection algorithms from the autonomous science literature. We evaluate each algorithm with respect to several autonomous geology applications. Tests show the algorithms’ performance on Mars Exploration Rover imagery, terrestrial images from analog environments, and synthetic images from a Mars terrain simulator. This provides insight into the detectors’ performance under different imaging conditions.

Detector	Method	Source
Rockfinder	Edge contours	[4]
MVJ (Modified Viola/Jones)	Filter Cascade	Carnegie Mellon
Rockster	Edge contours	[5]
Stereo	Stereo geometry	[8],[12]
Shadow (Marsokhod)	Shadow terminators	[9]
SQUID	Intensity Blobs	JPL
SVM	Classify pixels	JPL

Figure 1. Strategies for rock detection in images.

TABLE OF CONTENTS

1	INTRODUCTION	1
2	IMAGE DATASETS	1
3	TESTED ALGORITHMS	2
4	PERFORMANCE EVALUATION.....	4
5	EXPERIMENTAL RESULTS	5
6	DISCUSSION	7

1. INTRODUCTION

Detecting rocks in images is a valuable capability for autonomous planetary science. Rocks are pristine targets for compositional analysis with spectrometers. Their shape, size, and texture hold a wealth of geologic information. Computing the locations and distributions of rocks facilitates autonomous rover functions like adaptive target selection [4], selective image return [2], and autonomous site characterization [12]. Moreover, automatic rock detection can assist off-board data analysis. Rock size and shape distributions carry important geologic cues, but the manual analysis required to characterize these distributions is extremely time intensive. Currently scientists evaluate these statistics over very limited regions. Accurate rock detection would generate these rich data products quickly and cheaply.

Unfortunately rock detection is a difficult pattern recognition problem. Rocks exhibit diverse morphologies, colors and textures. They are often covered in dust, grouped into self-

occluding piles or partially embedded in terrain. Predictably there is no “silver bullet” algorithm to provide perfect accuracy in all these circumstances. However, one can still seek expedient rock detectors with performance that is sufficient for the desired application.

Research in the past decade has produced a variety of rock detection strategies (figure 1). These include stereo-based techniques for finding rocks based on their protrusion from the ground plane, edge-based methods that find closed contours, template-based methods that look for characteristic pixel patterns, and methods that detect rocks using their shadows. Each approach has advantages for different conditions and mission requirements. Nevertheless there has not been a formal comparison of these algorithms.

This work surveys a representative selection of rock detectors and evaluates their performance on four image datasets. We consider Mars Exploration Rover images from Pancam and Navcam instruments, along with images of physical laboratory analogs and computer-simulated terrain. These tests provide insight into both the algorithms’ comparative performance and the visual fidelity of laboratory analogues to actual Mars terrain.

2. IMAGE DATASETS

Images from four datasets (figure 2) were labeled manually to identify the ground-truth locations of all rocks. For each class of images we reserved a portion for training and tuning the algorithms. This section describes each of the datasets in greater detail.

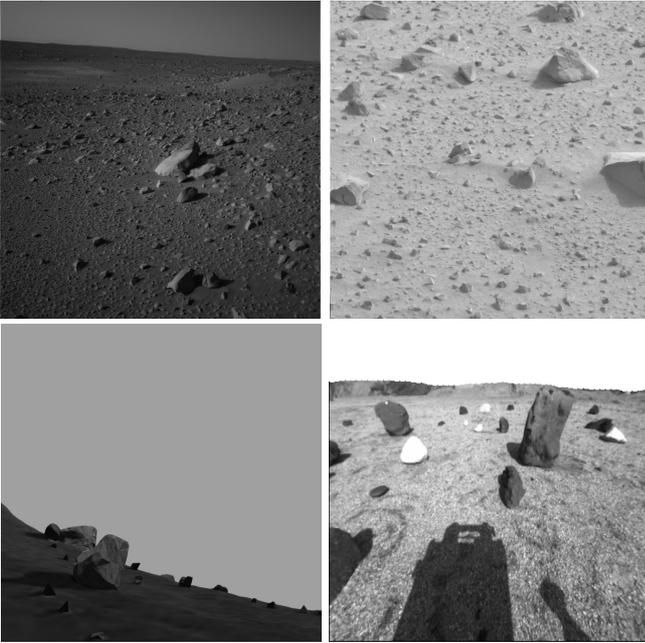


Figure 2. Images from the four datasets used in the performance comparison. Clockwise from upper left: MER Navcam, MER Pancam, Mars Yard, synthetic ROAMS.

MER Pancam Imagery

The first test set contained of 104 panoramic camera images taken by the Spirit Mars Exploration Rover (MER). We selected a range of mast elevation angles but favored image content within a 3-10m range of the rover. This range was deemed the most relevant for autonomous geology operations. Together the images comprised a dataset containing over 35,000 hand-labeled rocks. They also included a selection of typical non-rock content such as rover tracks, shadows and sediment features. We drew images from two sites: the “Mission Success” panorama from Spirit’s landing site and the “Legacy” panorama imaged during Sols 59 through 61. The images were acquired using Spirit’s 753nm bandpass filter.

Particle size analysis usually excludes the smallest rocks that are too numerous for accurate labelling. Similarly, we limited our analysis to those rocks greater than 4cm in length because they could be consistently identified across all datasets. We reserved an additional set of 12 Pancam images from Spirit rover imagery to serve as training data. These contained several thousand rocks from different locations along the Spirit mission path; like the test set they exhibited a range of terrain and lighting conditions.

MER Navcam Imagery

In addition to the panoramic camera imagery we considered 13 MER navigation camera images from the Spirit rover. These images contained over 13,000 hand-labeled rocks. We used Navcam images from three sites: initial im-

ages from the landing site, a panorama at sol 50, and a second panorama from sol 118. These monochrome images provided a wider field of view including both near-field objects and the horizon. They also showed a wider range of lighting conditions than the panoramic imagery; several were taken late in the Martian day with low-angle illumination and significant cast shadows.

Because relatively few labeled Navcam images were available we did not reserve any as a training set. Instead, we used the Pancam training set mentioned above for all tests involving Mars rover images. This compromise worked well for the specific detectors and datasets we considered. Again, we ignored rocks smaller than 4cm in size.

“Mars Yard” Analog Imagery

The “Mars Yard” is an outdoor rover testing environment at the NASA Jet Propulsion Laboratory (JPL) constructed to simulate Mars conditions. We used a sequence of Mars Yard images collected with the FIDO rover platform [4]. The Mars Yard image set consisted of 35 images and several hundred rocks. A second group of 28 Mars Yard images served as a training set.

The datasets were recorded by the FIDO rover’s mast-mounted navigation cameras during several rover operations tests. Image contents showed the sandy testbed terrain and a field of rocks that had been assembled for the tests. Background clutter like buildings and fences was removed from the images by hand. The rover shadow occasionally appeared in the images.

Synthetic Imagery

A fourth dataset utilized synthetic images of Mars terrain generated through the ROAMS rover simulator suite [10]. The ROAMS terrain simulator presented a Mars-like environment populated by rocks of various sizes. The simulated environment featured an undulating textured terrain, a realistic lighting model with cast shadows, and rocks with complex polygonal shapes. The ROAMS training and test sets contained 33 and 22 images respectively. Each set contained several hundred rocks.

3. TESTED ALGORITHMS

We drew representative algorithms from published studies and current research at Carnegie Mellon and the Jet Propulsion Laboratory (JPL). Each algorithm constitutes a complex system demanding a variety of unique design choices for pre-processing and parameter adjustment. Therefore, we caution against the presumption that any of these particular implementations achieved “optimum performance” for its technique. Nevertheless, we attempted to produce the best possible result from each detection system and used designers’ original source code whenever possible. Figure 3 shows the algorithms’ detection results on a common Navcam image. A brief description of each algorithm follows.

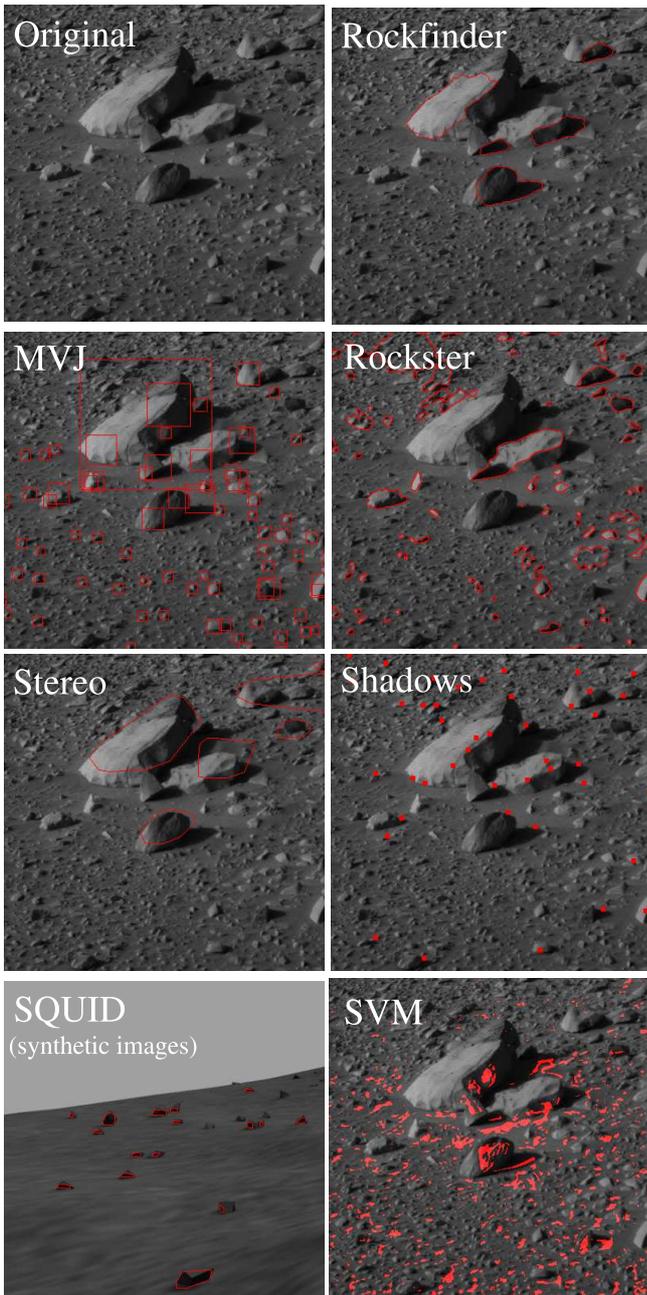


Figure 3. Red outlines show each algorithm’s detection results. Note that the SVM method does not actually find rocks at all, but merely labels individual pixels. The SQUID detection method at lower left only provides meaningful results on synthetic images.

Rockfinder

Rockfinder is an edge-based rock detection algorithm developed by the OASIS autonomous science group at JPL [4]. This detection technique uses a series of denoising and bilateral filter operations to remove spurious texture. Then it applies an edge detector (the Sobel or Canny algorithm) to find edge fragments and traces the resulting pieces into closed

contours. These closed shapes are unlikely to appear coincidentally in the Mars environment — the vast majority of closed edge shapes lie on the surface of rocks or are rocks themselves. The algorithm runs at multiple levels of a multi-scale image pyramid in order to detect rocks of all sizes.

Multiple Viola-Jones (MVJ)

The MVJ detector developed at Carnegie Mellon University exemplifies a template-based rock detection method. The algorithm utilizes the template cascade developed by Viola and Jones [13]. Here an Adaboost supervised training scheme builds a cascade of filters to quickly identify template windows containing rocks. This does not find the actual contour of the rock but it does provide an approximate bounding box that can then be used for spectrometer target selection or additional image processing (figure 3).

This rock detector uses a slight change of the classic Viola/Jones strategy; rather than train a single filter cascade we use multiple cascades to recognize rocks under different lighting directions. In principle one could train any number of independent detector cascades, but our version has only two detectors trained on rocks lit from the left and the right.

At runtime the multiple cascades compete to interpret the scene. The system detects rocks twice — once using each cascade — and uses only the results from the cascade that returns the greatest number of detections. The cascades learn to identify the characteristic shading patterns of a rock lit from a particular direction. These shading patterns seldom occur by accident so detection cascades associated with incorrect lighting conditions return few detections. The multiple cascade strategy identifies an approximate lighting direction even when the sun angle is not known with certainty.

Rockster

The “Rockster” algorithm is an edge-based algorithm developed at JPL [5]. While originally developed for synthetic images it has also shown promise for other datasets. The algorithm first identifies and removes any sky that is present at the top of the image. It then detects edges using an edge detector similar to the Canny algorithm. The next stage — an innovation of the Rockster algorithm — is a sequence of edge cleaning operations used to identify closed contours. These include breaking edges at points of high curvature, connecting endpoints with nearby edges to form T-junctions, and a gap-filling step that bridges missing contour segments. The final result is a set of closed shapes that can be extracted using a flood-fill technique. Rockster was successfully demonstrated on the FIDO platform in Sept. 2007 when it detected rocks during an integrated Mars Yard test of autonomous science operations.

Stereo Height Segmentation

Stereo range data is a component in several rock detectors including those by Gor [8], Fox [6] and Thompson [12].

These algorithms follow a common formula. They fit a planar ground model to the terrain using least-squares regression or RANSAC [7]. Then they find each pixel’s distance to the plane in order to create a height map. A segmentation of this height map identifies image regions that protrude above the surface.

We implemented a version of stereo hight segmentation for this work. Our version follows Gor’s strategy of detecting height-map discontinuities that indicate the tops of rocks. We apply a vertical derivative filter that responds strongly to these discontinuities. After finding the topmost rock pixels, a region-growing operation grows the rock regions down to a minimum height. The result is a segmentation of the rocks in the scene. Unlike the original Gor algorithm, we also apply a high-pass filter to the initial hieghtmap. This removes low-frequency changes in height like terrain bumps to improve performance when the ground is not perfectly planar.

Marsokhod Shadow Detector

We examine the rock detection algorithm used for image analysis during the 1999 Marsokhod rover field tests [9]. This method is one of the first examples of rock detection for autonomous science applications. It reduces the rock detection problem to one of finding shadows. Given a known sun angle, shadows suggest the location of rocks in the scene. The Marsokhod detector utilizes a spherical lighting model to predict the orientation of terminator lines (separating illuminated from darkened sides of the rock). It uses an edge detector to find candidate terminators, and selects appropriate candidates using edge orientations. While the result does not find the contour outline of the rock, it does identify a point on the rock’s surface that can be used as a target for spectroscopy.

SQUID

The *Smoothed Quick Uniform Intensity Detector* is a simple algorithm developed for use on synthetic images. First a bilateral filtering algorithm removes terrain texture. Then a simple intensity segmentation scheme identifies contiguous areas of constant pixel intensity. In this manner it locates contiguous blobs of intensity; a size filter discounts large blobs as belonging to the sky or terrain. Any remaining blobs are presumed to be rocks.

SVM Pixel Classification

Finally we investigate an algorithm that classifies every image pixel individually. These pixel classifications do not reveal the contours of each individual rock, but they can estimate the fractional coverage of rocks in the image. The algorithm uses a support vector machine classifier to analyze local windows around each pixel. It characterizes each pixel using a feature vector constructed from local intensity values. The SVM identified individual pixels that were most likely to lie on rocks.

4. PERFORMANCE EVALUATION

This section describes the metrics used to evaluate each rock detection algorithm. Pattern recognition systems are often judged according to a precision/recall profile: *precision* refers to the fraction of detections that are actually true rocks, while *recall* describes the fraction of rocks in the scene that are detected. In general the algorithms in these experiments do not provide a measure of detection confidence that would permit an explicit precision/recall tradeoff. However, the algorithms usually exhibit an obvious performance peak that strongly suggests a particular parameter setting.

We use several performance measures to evaluate test set performance for different autonomous geology tasks. First we evaluate a detector’s ability to perform target selection for applications like autonomous spectroscopy. This is tantamount to finding pixels that are most likely to lie on the surface of a rock. For algorithms that provided an outline or bounding box we used the centroid of the detected region as the target estimate. The relevant performance characteristic is accuracy - the fraction of chosen target pixels that actually lie on rock regions.

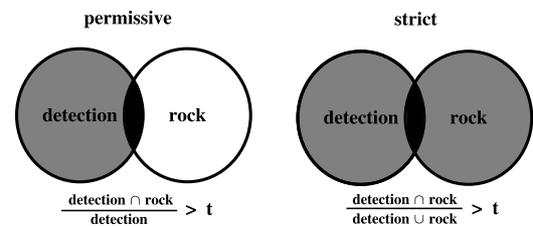


Figure 4. Graphical depiction of permissive (left) and strict (right) correspondence criteria. The ratio defined by the areas of intersection between detected contours and ground-truth rocks must be greater than a fixed threshold t . Both correspondence criteria permit at most one detection per rock.

Next we evaluate each detector’s detection precision and recall. To compute this score we find the area of overlap between detected and ground-truth rocks, and greedily associate the best-matching detections with their corresponding rocks. We match each ground-truth rock with at most one detection; unmatched rocks are labeled as false positives.

Each valid match must satisfy a correspondence criterion to ensure similarity between the detection and its associated ground truth rock. Performance scores are highly dependent on this correspondence standard so we evaluate two possibilities: a *permissive criterion* and a *strict criterion* (figure 4). The permissive criterion requires that more than 50% of the detected region contain the matched ground-truth rock. This guarantees that a majority of pixels in the detected region lie on the rock. However, it is forgiving of situations where the detector finds only part of the target. We also investigate a strict criterion that requires a 50% overlap between the regions’ intersection and their union. The strict criterion demands a close correspondence between the area of the true

rock and the detected region, so it better indicates the accuracy of automatically-computed rock attributes like size or shape.

Finally we evaluate autonomous prediction of the fractional area of terrain covered by rocks. We count detected rock pixels individually and compare the resulting count to the ground-truth fraction for each image. This standard does not require any correspondence between detected and actual rock pixels; it only concerns the total predicted number of rock pixels.

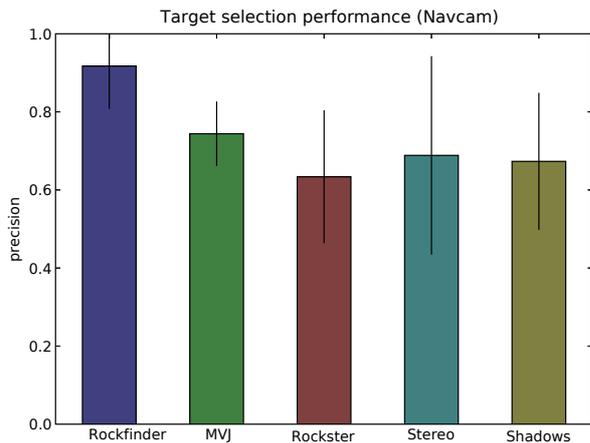


Figure 5. Target selection performance for Navcam imagery. The vertical axis represents the fraction of detected centroids that lie on a rock pixel.

5. EXPERIMENTAL RESULTS

Performance for Navcam and Pancam images is similar for all algorithms and tests with one exception — a range data misregistration error prevented us from evaluating the stereo algorithm on the Pancam dataset. Therefore the explanation that follows favors Navcam imagery except in cases where the Pancam dataset’s additional data is particularly informative.

Figure 5 shows target selection performance for navigation imagery. The Rockfinder algorithm has the best overall target selection performance; it reliably achieves accuracies over 90%. Failure modes differ for each detection algorithm. The Rockfinder algorithm occasionally mistakes long cast shadows for rocks (figure 6). The MVJ algorithm performs poorly in lighting conditions that differ from its training set (i.e. illumination from directly behind or in front of the camera). The stereo algorithm underperforms for small rocks or uneven terrain. Nevertheless, all rock detectors routinely score accuracies of 60 – 70% on the target selection task. In fact, these values may understate the performance of algorithms like Rockster and MVJ whose profuse detections could be filtered using size or shape heuristics.

Figure 7 shows average performance on Navcam images for rocks in four size categories. The permissive detection cri-

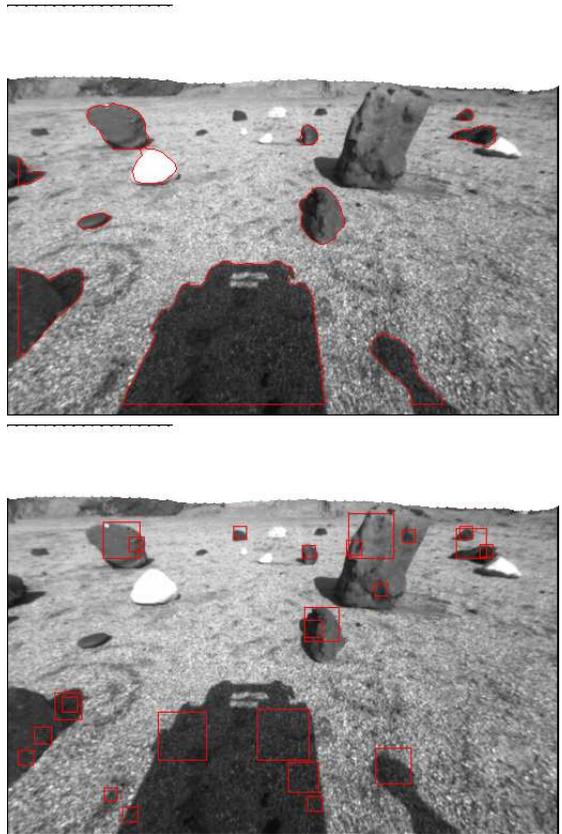


Figure 6. Rover shadows in the Mars Yard sequence confused most detection methods. Here Rockfinder (top) and the MVJ detector (bottom) both find one or more rocks in the rover shadow.

terion was used. Note that the MVJ algorithm’s performance increases with rock size, while the reverse is true for Rockster. Rockfinder and the stereo algorithm both perform poorly on rocks less than 10cm in length. This is expected in the case of stereo that relies on protrusion above the groundplane. However, there is no obvious mechanism that would cause a size bias in an edge-based method like Rockfinder. This algorithm regularly achieves 90% precision on medium to large rocks.

Figure 8 shows performance on Navcam images plotted by range categories. Most detectors show improved detection precision in the areas near to the rover. Not only do distant rocks subtend fewer pixels, but they also cluster more closely due to camera foreshortening.

Figure 9 shows the same detection results for Navcam imagery on an image-by-image basis. This provides some intuition about the variance that can be expected in performance from one image to the next. Again, both Stereo and Rockfinder algorithms exhibit high precision and low recall while the Rockfinder algorithm is the overall winner in precision. The Rockster algorithm achieves the highest overall recall for this test set at the cost of considerable inter-image variance.

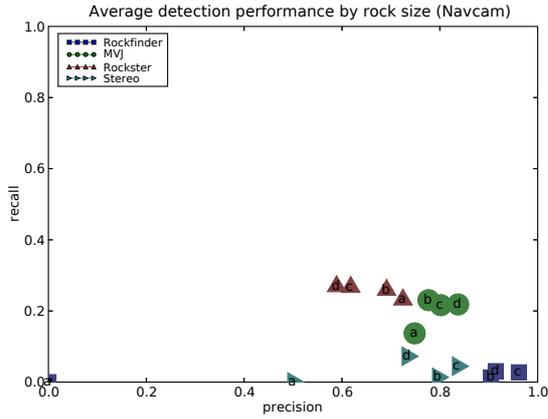


Figure 7. Performance by size for Navigation imagery. Points represent the average detection performance over the entire dataset for rocks in each size bin: a) $< 10cm$, b) $10 - 20cm$, c) $20 - 30cm$, d) $> 30cm$.

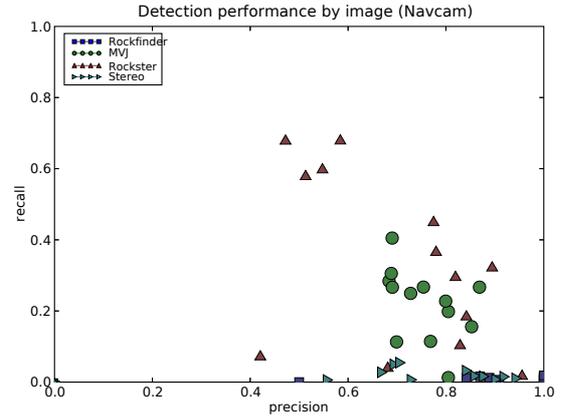


Figure 9. Rock detection performance for Navcam imagery (permissive criterion). Points represent the precision/recall result for each individual image

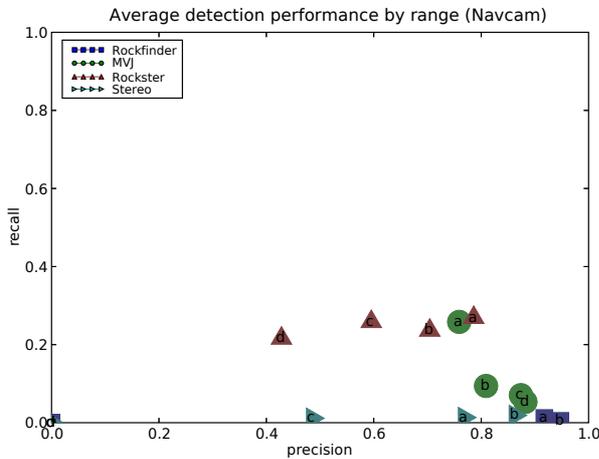


Figure 8. Performance by range for Navigation imagery. Points represent the average detection performance over the entire dataset for rocks in each range bin: a) $< 2m$, b) $2-4m$, c) $4-6m$, d) $> 6m$.

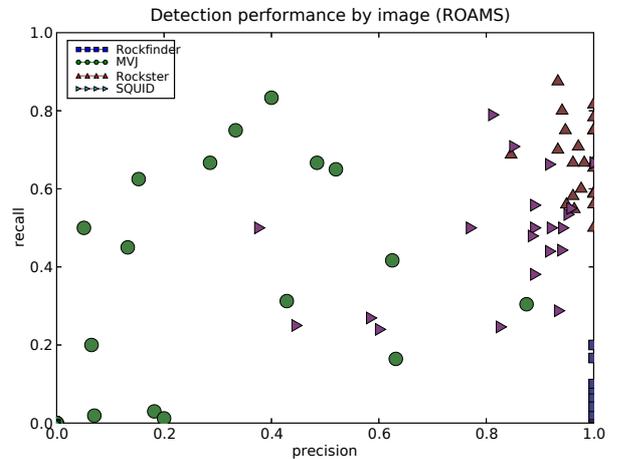


Figure 10. Rock detection performance for ROAMS imagery. The edge-based detection methods outperform MVJ for synthetic images.

Detection results for the other three datasets appear in figures 10, 11, and 12. Performance on the Mars-analog datasets is superior to performance on the actual Mars images. In particular, the Rockster and SQUID algorithms (that had been designed for use with ROAMS imagery) achieve much higher scores for the synthetic data. The obvious difference between Rockster’s performance in ROAMS and Navcam scenarios emphasizes the difference between the detection tasks for synthetic and real images (figures 9 and 10).

Recall on Mars Yard images is also higher than for Mars rover images. One might anticipate that rocks are easier to find in the simplified Mars Yard environment. The Mars Yard’s rocks are comparatively large and high-contrast. They are placed on the surface of the terrain by hand. Interestingly,

however, both MVJ and Rockfinder algorithms lose some precision on the Mars Yard imagery. This appears related to the rover shadow that appears in many Mars Yard test images (figure 6). Both MVJ and Rockfinder algorithms consistently identify this shadow as one or multiple rocks. In general, cast shadows are a significant issue for every algorithm except the Stereo method. They also impact detection performance on Mars imagery, where rocks’ shadows are often mistaken for rocks themselves.

Figure 13 shows Navcam detection results using the strict correspondence criterion. A comparison with figure 9 suggests that few detections actually capture the target rock’s true contour. The MVJ algorithm only detects bounding boxes and not contours, so its performance score suffers most from the stricter correspondence standard.

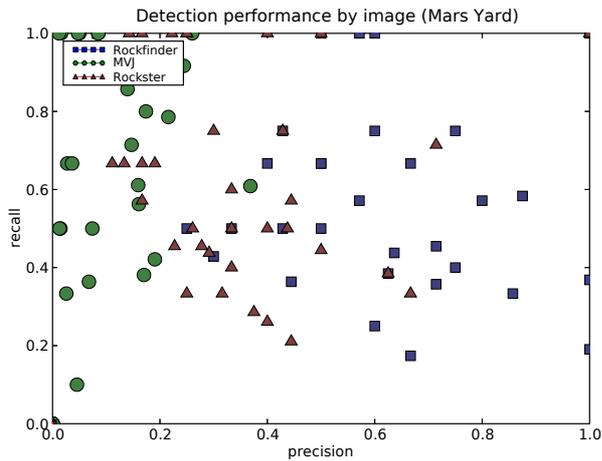


Figure 11. Rock detection performance for Mars Yard imagery. Each image contains just handful of rocks so occasional perfect scores appear.

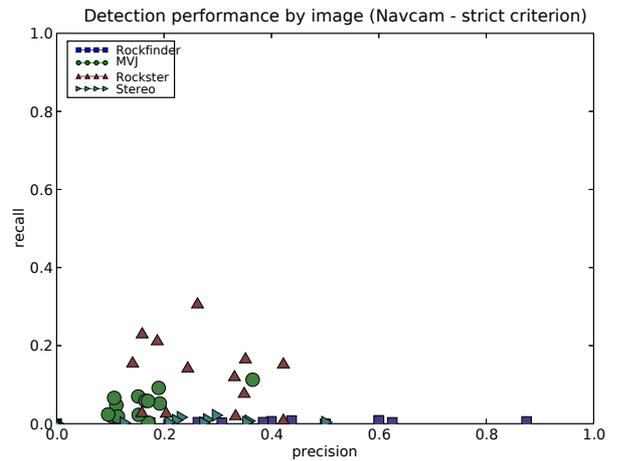


Figure 13. Rock detection performance for nav imagery (strict criterion).

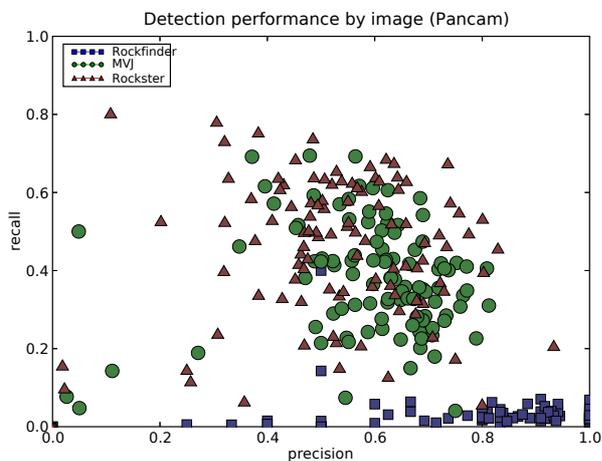


Figure 12. Rock detection performance for Pancam imagery. Performance is generally comparable to that for Navcam images, with MVJ and Rockster trading precision for recall.

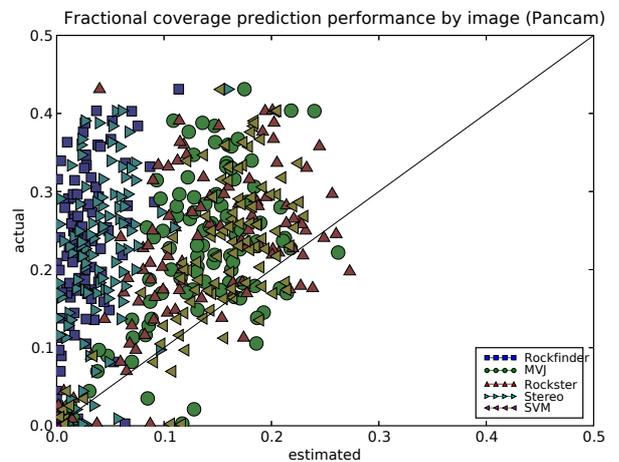


Figure 14. Coverage estimation performance for Pancam imagery. Each point corresponds to a single image. The horizontal axis represents the estimated fractional coverage of rocks, while the vertical shows the true value — a perfect estimate would lie on the diagonal line.

Finally, coverage estimation performance appears in figure 14). This task proves significantly more difficult — only the Rockster, SVM and Viola/Jones algorithms’ coverage estimates have any significant correlation with the ground truth value. Even these correlations are tenuous; at best they could make very coarse relative distinctions. The SVM method exhibits slightly lower variance than the other options. This is not surprising as the SVM algorithm was originally designed for the coverage estimation task.

6. DISCUSSION

The experiments of this work suggest some preliminary conclusions about the capabilities of current rock detectors. Target selection proved the easiest task, with Rockfinder demonstrating accuracy rates of over 80%. Usage restrictions on

range or time-of-day (to reduce cast shadows) might reduce error even further. Note that rocks covered 0–40% of the terrain in the tested datasets. This suggests that even the worst-performing rock detectors could already yield significant improvement over blind target selection.

Finding all the rocks in an image is more difficult. Recall was characterized by poor performance and size biases. Even with the permissive correspondence criterion no detector averaged more than 60% recall on any field dataset. In other words, none of the detectors we evaluated can yet produce reliable statistics about the visual features of rocks in a Mars image. Research has suggested that an imperfect detector might still be able to make meaningful distinctions between neighbor-

ing locales [12]. Nevertheless, absolute quantitative site descriptions appear beyond the capability of current detectors. The coverage estimation task echoes this finding. Coverage estimation is presumably an easy site-characterization task with better ground-truth fidelity than subtle rock size or shape analysis. Nevertheless, the high variance of detectors' coverage estimates suggest that autonomous detection is still no replacement for traditional manual analysis.

This discrepancy between the ease of target selection and the difficulty of detection suggests that autonomous site characterization might be improved by pairing visual analysis with additional sensors. Rovers can complement a preliminary visual identification of rocks with follow-up spectroscopy or thermal imagery that is both geologically informative and easier to interpret autonomously [1].

A more sophisticated rock detector may still make reliable detection a reality. A recent resurgence in general object recognition research will certainly yield insight into the rock detection problem. It may also be possible to combine several detection algorithms for greater accuracy. Complementary algorithms running in parallel might improve performance either by finding a wider range of rocks or validating each other's detections. In the meantime, autonomous geology systems that focus on target selection for instrument deployment could play to the best strengths of rover autonomy.

ACKNOWLEDGMENTS

We are thankful for the assistance of Benjamin Bornstein, Michael Burl, Michele Judd, Daniel Gaines and the OASIS team. Thanks are also due to Ted Roush for the use of Marsokhod detection code.

REFERENCES

- [1] B. Bornstein and R. Castaño. "Creation and Testing of an Artificial Neural Network Based Carbonate Detector for Mars Rovers," *Proceedings of the IEEE Aerospace Conf.*, Big Sky, Montana, 2005.
- [2] 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," *Proceedings of the IEEE Aerospace Conf.*, Big Sky, Montana, 2003.
- [3] A. Castaño, R. C. Anderson, R. Castaño, T. Estlin, and M. Judd, "Intensity-based rock detection for acquiring onboard rover science," *Lunar and Planetary Science*, 35, 2004.
- [4] R. Castaño, M. Judd, T. Estlin, R. C. Anderson, D. Gaines, A. Castano, B. Bornstein, T. Stough, and K. Wagstaff. "Current Results from a Rover Science Data Analysis System", *Proceedings of the IEEE Aerospace Conference*, Big Sky, Montana, 2006.

- [5] R. Castaño, T. Estlin, D. Gaines, C. Chouinard, B. Bornstein, R. C. Anderson, M. Burl, D. Thompson, and M. Judd. "Onboard Autonomous Rover Science", *Proceedings of the IEEE Aerospace Conference*, Big Sky, Montana, 2007.
- [6] J. Fox, R. Castaño, R. C. Anderson, "Onboard autonomous rock shape analysis for Mars rovers," *Proceedings of the IEEE Aerospace Conference*, Big Sky, Montana, 2002.
- [7] M.A., Fischler and R.C Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography." *Communications of the ACM*, 24:6, 1981.
- [8] V. Gor, R. Castaño, R. Manduchi, R. C. Anderson, and E. Mjolsness, "Autonomous rock detection for Mars terrain," *Proceedings of AIAA Space 2001*, Albuquerque, August 2000.
- [9] V. C. Gulick, R. L. Morris, M. A. Ruzon, and T. L. Roush, "Autonomous image analysis during the 1999 Marsokhod rover field test," *J. Geophysical Research*, 106, No. E4, 7745–7764, 2001.
- [10] A. Jain, J. Balaram, J. Cameron, J. Guineau, C. Lim, M. Pomerantz, and G. Sohl. "Recent Developments in the ROAMS Planetary Rover Simulation Environment," *Proceedings of the IEEE Aerospace Conference*, Big Sky, Montana, 2004.
- [11] L. Pedersen, *Robotic Rock Classification and Autonomous Exploration*, PhD thesis, Robotics Institute, Carnegie Mellon University, CMU-RI-TR-01-14.
- [12] D. Thompson, T. Smith and D. Wettergreen, "Data Mining During Rover Traverse: From Images to Geologic Signatures," *ISAIRAS 2005*.
- [13] P. Viola and M. Jones, "Rapid object detection using a boosted cascade of simple features." *CVPR 2001*.
- [14] M. D. Wagner, D. Apostolopoulos, K. Shillcutt, B. Shamah, R. G. Simmons, W. Whittaker, "The Science Autonomy System of the Nomad Robot," *ICRA 2001*, 2, 1742—1749.
- [15] K. L. Wagstaff, R. Castaño, S. Dolinar, M. Klimesh, R. Mukai, "Science-based region-of-interest image compression," *Lunar and Planetary Science*, 35, 2004.



David R. Thompson PhD student, Robotics Institute at Carnegie Mellon University. He received a MSc. in Informatics from the University of Edinburgh. His work has focused on computer vision and machine learning technologies for autonomous science. Recent projects include autonomous site characterization and geologic mapping during rover operations in the Atacama Desert of Chile.



Dr. Rebecca Castaño Supervisor, Machine Learning Systems group at JPL, OASIS Technical team lead. She received her Ph.D. in Electrical Engineering from the University of Illinois with her dissertation in the area of computer vision. Dr. Castaño has been advancing the state of the art in onboard science analysis methods for the past eight years and has been lead author on numerous publications in this field. Her research interests include machine learning, computer vision and pattern recognition.