

Techniques for Onboard Prioritization of Science Data for Transmission

R. Castaño,¹ R. C. Anderson,² T. Estlin,¹ D. DeCoste,¹ F. Fisher,¹
D. Gaines,¹ D. Mazzoni,¹ and M. Judd¹

Future planetary exploration missions will continue to face an ever-increasing data prioritization problem as instrument data collection rates continue to exceed spacecraft downlink transmission rates. Hard decisions must be made about what data are sent back to Earth and what data are purged without being seen by scientists. In this article, we present a suite of techniques for the prioritization of science data for transmission. These techniques include methods to ensure that data representative of the local geology, as well as any unusual observations, are given the highest priority.

I. Introduction

As planetary exploration continues to expand, the combination of more missions, an increased number of instruments, and the advanced capabilities of those instruments will cause an increase in the volume of data to be transmitted back to Earth via the Deep Space Network (DSN). Missions will have to make critical decisions regarding the quantity and quality of the downlinked data, both of which are indirectly affected by, among other things, the availability constraints of the DSN and the limited spacecraft power. Although the DSN's receiving capability increases every year, the number of missions it must service is growing rapidly.

New methods must be developed to maximize the science return for the available bandwidth. Conventional data compression helps in this regard; however, ever-increasing amounts of compression can cause unacceptably high distortion levels. We are developing onboard analysis methods to autonomously prioritize data for downlink as another approach to getting the most out of the limited bandwidth.

Data prioritization already implicitly occurs. A priori decisions are made regarding when to turn on an instrument based on the number of data sets that can be transmitted to Earth. Instruments may not be activated because existing scenarios have limited or no ability to make intelligent decisions about the value of data being collected. In contrast, by collecting more data than can be transmitted to Earth and applying onboard data-analysis technology to carefully select which data are sent, the quality of data returned can be increased, maximizing the science return for the available downlink. Selecting the most

¹ Exploration Systems Autonomy Section.

² Planetary and Life Detection Section.

The research described in this publication was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

interesting data from a set requires encapsulation of data characteristics with high science value in a form that can be analyzed and evaluated onboard. The spacecraft must have the capability to automatically recognize data that contain information about targets with high science interest.

Our operating scenario is rover traverse science, although the algorithms developed in this work can easily be adapted for use on orbital platforms. Rover traverse science involves the analysis of data that are collected as a rover travels between study sites. Since images are among the highest volume data acquired by the rover and thus consume the most bandwidth, images represent one of the greatest opportunities for data prioritization. Therefore, in this article, we focus on prioritization of images. To assess and subsequently prioritize the scientific value of a set of collected images, we first must extract the information found within the images. A geologist in the field would extract the information from a site by identifying geologic features, including the albedo, texture, shape, size, color, and arrangement of rocks, and features of the topography, such as layers in a cliff face. Based on knowledge and experience, the importance of these features then would be evaluated. Our system begins to emulate this process by first evaluating each image and locating the rocks within it. Next, we extract the properties of the rocks, including albedo, texture, and shape features. These features are assigned a science value, and the importance of the image is evaluated based on these values as compared with the features extracted from rocks identified in other images. Images with interesting features, such as rocks with unusual shapes or textures, should be ranked higher than images without distinctive features.

We have developed three different prioritization methods that use the extracted rock features to rank the rocks in terms of scientific importance. Once the rocks are prioritized, the images containing the rocks are ranked based on the rocks contained within them. The first technique, target signature detection, recognizes pre-specified signatures that have been identified by the science team as data of high interest. The second technique, novelty detection, identifies unusual signatures that do not conform to the statistical norm for the region. The last method, representative sampling, prioritizes data for downlink by ensuring that representative rocks of the traversed region are returned. These three data prioritization and analysis techniques have been identified by mission scientists as critical onboard rover traverse science data-analysis capabilities.

Without extensive ground testing and validation, the science team will be extremely reluctant to use autonomous, onboard prioritization of data for downlink. As a validation mechanism, we have developed a robust method for quantifying the correlation between our automated prioritization and a scientist's prioritization of the same data set.

Finally, prioritization can be used for more than just data downlink decisions. It also can be used to trigger opportunistic science observations. Prioritization that calls for opportunistic science cannot be utilized without a method of re-sequencing the rover, or orbital spacecraft, to obtain the additional scientific observations requested. This ability for real-time opportunistic science requires integrating the prioritization module with the onboard planning and scheduling system. We briefly outline our approach to this integration.

Throughout the remainder of this article, we describe our methods for feature extraction and prioritization. We then briefly outline our validation approach and discuss future directions for the project.

II. Feature Extraction

The first step in image prioritization is to extract geological features of interest from the image itself. Our work has focused on the analysis of rocks in the scene, and thus we begin by locating rocks in a stereo image pair. Previous methods for locating rocks in an image include techniques using shadows and information about the Sun angle [11]. Our technique for locating rocks is based on finding objects above the ground plane. We begin by determining the ground plane from the stereo range data. We then produce a height image, in which the value of each pixel represents the elevation of the point above the

ground plane. Level contours in the height image are calculated and then these contours are connected from peaks to the ground plane to identify the rocks [10]. Rock properties including albedo, visual texture, and shape then are extracted from the identified rocks.

The reflectance properties of a rock provide information about its mineralogical composition. We measure albedo, an indicator of the reflectance properties of a surface, by computing the average gray-scale value of the pixels that comprise the image of the rock. Shadows and Sun angle can affect the gray-scale value of a pixel. Although this can be corrected by using the range data along with knowledge of both the Sun angle and the camera orientation, our foundational work discussed in this article does not address these specific issues.

The second rock property extracted is visual texture. Visual texture can provide valuable clues to both the mineral composition and geological history of a rock (see Fig. 1). Visual texture can be described by gray-scale intensity variations at different orientations and spatial frequencies within the image. We measure texture using a bank of Gabor filters [9]. A Gabor filter is a complex exponential modulated by a Gaussian. Taking the Fourier transform, the filter is given by a Gaussian in the frequency domain,

$$G(u, v, \theta, F, \sigma_u, \sigma_v) = e^{-2\pi^2([(u'-F)/\sigma_u]^2 + [(v')^2/\sigma_v^2])}$$

where u and v are the horizontal and vertical frequency coordinates, u' and v' are the rotated frequency coordinates ($u' = u \cos \theta + v \sin \theta$ and $v' = u \sin \theta + v \cos \theta$), θ is the angle of orientation, F is the center frequency or scale, and σ_u and σ_v define the bandwidth of the filter. The benefit of these filters for texture is that they are local spatial bandpass filters. Moreover, the filters achieve the theoretical limit for conjoint resolution of information in the 2-D spatial and 2-D Fourier domains [5]. Thus, in a sense, they are the optimal filter for measuring the frequency response within a local region of the image. The conditions under which a class of 2-D Gabor wavelets can provide a complete representation of an image can be identified [12]. In this work, we want a set of filters that can be used to describe the texture of regions of an image (i.e., the rocks), but we do not need to reconstruct the image. The family of self-similar filters that we use is defined by the minimum and maximum frequencies passed by the filter set, as well as the number of orientations (linearly spaced) and the number of scales (logarithmically spaced), where we assume the aspect ratio of the Gaussians is 2:1 and the filters intersect radially at the half-amplitude. An example of such a set of filters in the frequency domain is shown in Fig. 2. The texture response for an image is measured by convolving each filter in the set with the original image. We then take the magnitude of the response at a pixel. The resulting image cube has one layer corresponding to each filter, i.e., at each pixel there is a real-valued vector representing the measured texture centered on the pixel.




Fig. 1. Examples of visual texture.




Fig. 2. The magnitude of frequency response for a bank of 12 Gabor filters with 3 scales and 4 orientations (0 Hz is at 125 pixels).

Due to onboard computational constraints, a compromise must be made between the number of filters used and the discriminatory power of the filter bank. With the addition of filters to the set (i.e., more scales or orientations), each filter covers a smaller region of the frequency domain, enabling distinction of finer-frequency signals and providing more discriminating power between different textures in the spatial domain. We have found that 12 filters (four orientations and three scales) are effective without being computationally prohibitive [2].

Our current system is not rotationally invariant; however, by merging orientation information across each scale, the sensitivity to variation in the rotation of rocks within the image could be eliminated. This is primarily a consideration for rocks whose geological texture has an orientation preference such as layering or the orientation of grains that can be present in metamorphic rocks.

Another important and geologically useful feature of rocks is their shape. For example, a rock that is highly rounded may have undergone fluvial processing and traveled far from its source. Conversely, a rock that is highly angular is likely to be close to its source and to have undergone minimal secondary processing. The shape of a rock can be difficult to describe precisely. Our system describes the shape of a rock in an image using three parameters that capture how close to circular and how angular the rock is [8]. We begin by fitting an ellipse to the boundary points of the identified rock in the image. Our first shape measure is the eccentricity of this ellipse. Our second measure is the estimated mean-squared error between the boundary points and the ellipse. The third measure is angularity. Our angularity measure is based on an algorithm developed by Chetverikov et al. [3] to detect corners of high angularity in images. Assuming the boundary points of the rock are labeled in the clockwise direction starting from an arbitrary point, we wish to estimate the curvature at a particular point, P_i . We define the estimated curvature at point P_i as

$$\rho_i = \max_{\substack{h,j \\ h=i-\omega, \omega \in [\alpha, \dots, \beta] \\ j=i+\omega, \omega \in [\alpha, \dots, \beta]}} \angle P_h P_i P_j$$

where $\angle P_h P_i P_j$ denotes the angle between line segments $\overline{P_h P_i}$ and $\overline{P_i P_j}$, and α and β are two predefined integer constants. The indices h and i are computed modulo the number of boundary points so that they

wrap around the boundary. In earlier work, $\alpha = 2$ and $\beta = 9$ were determined to be the most effective values for this application [8].

Having calculated an estimated curvature, ρ_i , at every point on the boundary, we define the angularity of a rock as the standard deviation of the estimated curvature at each point. Thus, for a rock with N boundary points,

$$A_{2d} = \sqrt{\frac{1}{N-1} \sum_{i=1}^N (\rho_i - \bar{\rho})^2}$$

where $\bar{\rho}$ is the average of the curvature of the N points. Using this measure, a rock with a perfectly round boundary would have zero angularity.

Image analysis is two-dimensional; however, rocks are three-dimensional objects. We currently are incorporating three-dimensional properties into our system. These include estimates of sphericity and 3-D angularity, both of which can be derived from stereo range data.

III. Prioritization of Rock Image Data

Once the rocks have been identified in each image, the features of each rock (e.g., shape, albedo, visual texture) are extracted. The extracted values of the features are concatenated to form a feature vector, representing the quantified properties of the rock. The feature vectors for each rock are input into the three distinct prioritization algorithms described in this section.

Each of the different prioritization techniques provides a separate ranking of the images, which in turn is based on the most important rock found in each image. The examples in this section, presented to demonstrate the different prioritizations, use a data set containing 178 rocks extracted from 25 images. These images were taken by the Field Integrated Design and Operations (FIDO) rover, an experimental test rover for the 2003 Mars Exploration Rovers, during a field test near Flagstaff, Arizona, in August 2002.

A. Target Signature

A mission's instruments are carefully selected to collect information that will provide valuable insight about the history of, or current conditions on, the planet. Certain instruments, therefore, will be used to search for key target signatures that indicate the presence of what scientists consider crucial information. Thus, when only limited data can be sent to Earth, it is very important to scientists that any data containing these signatures be among the data that are returned.

Target signatures are specified by identifying nominal values for each of the relevant features. An importance is then assigned to each of the features. Rocks are prioritized as a function of the weighted Euclidean distance of their extracted feature vector from the specified feature vector.

We have implemented an efficient and easy to use graphical user interface (GUI) for scientists to stipulate the value and importance to assign to each feature. This GUI can be used in two ways. The first method involves manually specifying a set of feature vectors, as shown in Fig. 3. In the example, the scientist has chosen to prioritize rocks based on two aspects of their shape, eccentricity, and ellipse fit. This specification is used to rank the data set based on how round the rocks are, i.e., high priority to rocks with low eccentricity and low ellipse fit error. The second manner in which scientists can specify a target signature is by selecting a rock with interesting properties from the set of already identified rocks. Rocks similar to the selected rock with respect to the specified properties are given a high priority.




Fig. 3. Target signature specification and rock prioritization: (a) target signature GUI for selecting round rocks, i.e., low eccentricity and good fit to an ellipse, (b) image containing one of the best-fitting (round) rocks to the designated target signature, and (c) image containing one of the worst-fitting (not round) rocks to the specified target signature.

B. Novelty Detection

We have developed three methods for detecting and prioritizing novel rocks, representing the three dominant flavors of machine learning approaches to novelty detection: a distance-based method, a probability-based (i.e., “generative”) method, and a discriminative method. These methods for novelty detection are applicable to a variety of novelty-detection tasks but are specifically designed with onboard constraints and large candidate feature spaces in mind.

The first novelty detection method is a distance-based, K -means, clustering approach [1]. K -means is an unsupervised clustering algorithm. In unsupervised procedures, the clustering is performed using the data themselves, with no information provided as to what data belong to any of the classes. This is in contrast to supervised methods, in which a provided set of training data that is labeled to show category membership is used to design the classifier [6]. We currently employ K -means due to its relatively low computational requirements, although any unsupervised method could be used. Initially, the K -means clustering algorithm is applied to all available rock data. The K -means algorithm alternates between computing cluster centroids and assigning each feature vector to the cluster represented by the nearest centroid. The novelty of any rock is defined to be the distance from the rock feature vector to the nearest centroid of any of the k clusters.

Since the feature vector consists of disparate components, such as albedo and shape, consideration must be given as to how the vectors will be compared. A number of normalization options were considered, including normalizing each subvector, e.g., the two components associated with albedo would be normalized to a unit length and the twelve components representing texture would be normalized to a unit length separately. Distance then is measured by comparing the angles between albedo vectors, etc. At present, we are using the z -norm, where every element of the vector is normalized to a standard normal distribution over the sample set. We use Euclidean distance to compare normalized feature vectors, thus effectively employing a Mahalanobis distance to compare vectors. This normalizing method can be applied without contextual knowledge of the relationships between the components of the feature vector. We anticipate that, using input provided by scientists, the weighting and distance metric could be adjusted to return even more scientifically significant results.

The second technique uses a Gaussian mixture model of the probability density over the rock feature space to estimate the novelty of a rock. The density of the distribution from which the feature vectors are drawn is assumed to be a linear combination of component Gaussian densities, $p(\bar{x}) = \sum_{j=1}^M p(\bar{x}|j)P(j)$, where $p(\bar{x})$ is the probability of feature vector \bar{x} , $p(\bar{x}|j) \sim N(\theta_j, \sigma_j)$ is the probability of generating feature vector \bar{x} given that \bar{x} is generated from component j , $P(j)$ is the prior probability of a feature vector having been generated from component j of the mixture, and M is the number of Gaussian components in the mixture model. M is selected to maximize the expressiveness of the model (more Gaussians can accurately model complicated surfaces) while minimizing computational resource requirements. Rock feature vectors that have been previously collected are used as training data in the Expectation-Maximization algorithm [1] to estimate the values of θ_j , σ_j , and $P(j)$ for $j = 1 \cdots M$. The novelty of a rock is inversely proportional to the probability of that rock being generated by the learned model.

The final method is a discrimination-based kernel one-class classifier approach. Here we treat all previous rock data as the “positive class” and learn the discriminant boundary that encloses all that data in the feature space. We essentially consider the previous rock data as a cloud scatter in some D -dimensional space, where D is the number of features. The algorithm learns the boundary of that cloud, so that future rock data that fall outside the cloud boundary are considered novel. The further the new rock feature is from the boundary, the more novel it is.

An example of using the third algorithm on the test data set of rocks extracted from the FIDO images is presented in Fig. 4. Applying a direct mapping of the rock prioritizations to prioritize image data for downlink would give the image containing the most novel rock the highest priority. Out of the 25-image



Fig. 4. Detection of a significant novel rock. One of the most novel rocks detected in the test data set is the rock indicated in this image. The marked rock is a piece of petrified wood that is extremely interesting to scientists. Unfortunately, during the course of the FIDO field test, the petrified wood was not identified by the remotely located geologists.

data set, the image in Fig. 4 was ranked sixth highest in terms of novelty and, thus, would be included for downlink under a scenario that transmits six or more of the most novel images. The marked rock is a piece of petrified wood that is extremely interesting to scientists. Of the entire set, it has been designated as one of the most important rocks by scientists on our team who both studied the image set and were at the field site. In contrast, if six images were randomly selected from the set of 25, there is only a 24 percent chance that this image would be included in the downlink set. In the future, the algorithm will need to be tuned with input from scientists to identify rocks that are scientifically novel; however, initial tests, including the example shown, already have shown promise without yet using domain information.

C. Representative Target

One of the objectives for rover traverse science is to gain an understanding of the region being traversed. To meet this objective, the data downlinked should include information on rocks that are typical for a region, not just information on interesting and unusual rocks. A region is likely populated by several types of rocks with different relative abundances. If uniform sampling is employed for downlink image selection, as opposed to our autonomous onboard selection process, the downlinked set will be biased towards the dominant class of rock present. This situation may result in less populated classes (i.e., less common rocks) not being represented at all in the downlinked data.

In our representative sampling algorithm, the rocks are clustered into groups based on their feature vectors using K -means. The science team pre-assigns a weighting, or importance value, to each property (shape, albedo, texture) in the feature vector. Different weight assignments can be used to emphasize the properties of highest interest. For example, albedo and texture typically are used to distinguish types of rocks, but rock size may be used if sorting is of interest. (Rocks that have been subjected to a geologic process such as flooding may be sorted over a surface area according to their size. The presence of such sorting provides information about the processes that may have occurred.)

The data then are prioritized to ensure that representative rocks from each class are sampled. For each class of rocks, we find the most representative rock in the class, i.e., the single rock in any image

that is closest to the mean of the set. We give a high priority to the image containing this rock. The process is complementary to novelty detection using K -means, where rocks that are farthest from the cluster means are given the highest priority. Rocks within each class are ranked in order of increasing distance from cluster centers, and rocks (in different sets) with the same rank are given equal priority in overall ranking. If a well-ordered composite list is required, rocks of the same rank can be sorted based on the distance to their respective cluster centers.

Currently, the value of K , the number of classes, is fixed prior to performing prioritization. There are at least two methods that we could use to determine a suitable value for K . First, the number of classes to use could be determined using the data in an unsupervised manner, e.g., via cross-validation techniques [14]. The second method of determining the number of classes would be to use manual intervention. After the rover sends down initial images, the number of classes could be determined by experts on the ground. This value then would be uplinked and the onboard system updated.

An example of prioritization based on representative rocks using the August FIDO image set is shown in Figs. 5 and 6, where we have defined three classes of rocks, labeled A, B, and C. The top six rocks, the top two from each of the three classes, are shown in Fig. 5. Thus, if it were possible to transmit only three of the images from the entire set, transmitting the three images shown in Fig. 6 ensures that an example of each of the classes reaches the scientists. If six images are sent, the downlink set would include two representative rocks from each class.

In the future, the spatial locations of rocks will be used in addition to extracted visual properties to enable expanded analyses. This includes characterizing local surface regions, i.e., distinguishing between regions that have distinct distributions of abundances and types of rocks, and sorting, which requires size and location information.

IV. Validating Prioritization Algorithms

In order for our autonomous prioritization algorithms to be accepted and used, we need a technique for validating the results of our methods. In particular, we would like a quantitative measure to gauge how closely our algorithms match the priorities of experts. Although we will not present detailed results here, we briefly describe our approach to validating the prioritization techniques.

Our validation approach is to gather sample prioritizations from expert planetary geologists on several collections of images. We have implemented a Web-based application GUI that enables experts to prioritize images and add annotations for their decisions. Our Web-based expert prioritization application allows planetary geologists from around the world to access collections of images and to provide their input into how the images should be prioritized. Figure 7 shows an example of the results from an individual expert ranking on a set of five images.

After collecting information from the experts, statistical methods are used to combine the results from a number of experts to compare consistency across the experts and to compare their results with the prioritizations produced by our algorithms [13]. This process, which is based on accepted statistical methods for combining and comparing rankings, provides a quantitative measure of the performance of our algorithms.

V. Planning and Scheduling

In addition to downlink data selection, the prioritized rock information can be used for opportunistic science. Based on the rock prioritizations, a new set of targets for collecting additional data can be formed. These targets can be passed to other onboard autonomy software that will modify the onboard command sequence in order to collect the new science data. Current approaches to rover control require



Fig. 5. Prioritized ranking of rocks for three clusters of rocks. Using the representative target prioritization technique ensures that images containing the most representative rock of each class have the highest priority for downlink.





Fig. 6. The three most important images based on the representative rock prioritization shown in Fig. 5: (a) Class A, (b) Class B, and (c) Class C.

Priorities from Your Information






Rank	Image	Score
1	Image 2 	4
2	Image 1 	3
3	Image 3 	2
4	Image 0 	1
5	Image 4 	0

Fig. 7. Example results for an expert prioritization of a set of five images.

human analysis to determine goals and manually convert the set of high-level science goals into low-level rover command sequences. By integrating these components onboard, we enable a rover to collect data when an unexpected but scientifically interesting target is encountered. The loop between planning, data analysis, and data collection is closed onboard the rover, as shown in Fig. 8, enabling new data to be collected with little or no communication with Earth.

The capability of changing the activities performed by a rover is provided by the Continuous Activity Scheduling Planning Execution and Replanning (CASPER) planning and scheduling system [4,7], which can

- (1) Autonomously evaluate whether goals to collect data for new science targets can be achieved given the state of the rover
- (2) Modify the current command sequence to incorporate new targets
- (3) Monitor execution of that sequence in case further adjustments are necessary

The CASPER planning and scheduling system evaluates an input set of goals, the rover's current state, and resource levels. For example, the goal might be to take another measurement. If the current state is that the spectrometer is pointing away from the intended target, there must be sufficient power available to point the spectrometer. CASPER generates a new sequence of commands that satisfies as many of the new goals as possible while obeying any rover resource and operation constraints. Goals are evaluated based on their priority (as assigned by the data analysis software). If limited resources are available, then only the highest priority goals may be included in the new plan (i.e., command sequence).

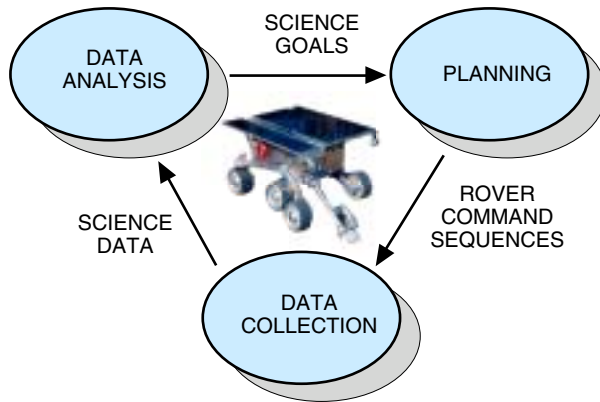


Fig. 8. Closed-loop data flow of the traverse science system.

VI. Conclusions

Although we have used rocks to demonstrate the benefits of autonomously prioritizing data, the prioritization techniques can be applied to any feature information that can be extracted from the data. Thus, they apply to more complex gray-scale image features, other data modalities (e.g., hyperspectral images), and orbital, as well as ground, data applications.

The DSN will remain a valuable, yet constrained, resource for future deep-space missions as the number of high-bandwidth missions increases. Traditional data compression can cope with only so much of this increase. To maximize science return, it is first necessary to autonomously identify the data with the highest science value, while the data are still onboard the spacecraft, and then to return the prioritized data to the science team. This requires a quantifiable measure of science value that can be evaluated onboard and a prioritization mechanism to rank the data for downlink based on the measured science value. We have described a system that implements this functionality by extracting properties from image data and prioritizing the data represented by the extracted properties using three distinct data prioritization methods that together allow a mission to achieve its primary scientific exploration objectives.

References

- [1] C. Bishop, *Neural Networks for Pattern Recognition*, New York: Oxford University Press, 1995.
- [2] R. L. Castaño, T. Mann, and E. Mjolsness, "Texture Analysis for Mars Rover Images," *Proceedings of Applications of Digital Image Processing XXII*, SPIE vol. 3808, pp. 162–173, July 1999.
- [3] D. Chetverikov and Z. Szabo, "A Simple and Efficient Algorithm for Detection of High Curvature Points in Planar Curves," *Proceedings of the 23rd Workshop of the Austrian Pattern Recognition Group*, Steyr, Austria, pp. 175–184, 1999.

- [4] S. Chien, R. Knight, S. Stechert, R. Sherwood, and G. Rabideau, “Using Iterative Repair to Improve Responsiveness of Planning and Scheduling,” *Proceedings of the 5th International Conference on Artificial Intelligence Planning and Scheduling*, Breckenridge, Colorado, pp. 300–307, April 2000.
- [5] J. G. Daugman, “Uncertainty Relation for Resolution in Space, Spatial Frequency, and Orientation Optimized by Two-Dimensional Visual Cortical Filters,” *Journal of the Optical Society of America*, vol. 2., no. 7, pp. 1160–1169, 1985.
- [6] R. Duda and P. Hart, *Pattern Classification and Scene Analysis*, New York: John Wiley and Sons, 1973.
- [7] T. Estlin, A. Gray, T. Mann, G. Rabideau, R. Castaño, S. Chien, and E. Mjolsness, “An Integrated System for Multi-Rover Scientific Exploration,” *Proceedings of the Sixteenth National Conference on Artificial Intelligence*, Orlando, Florida, pp. 613–620, August 1999.
- [8] J. Fox, R. Castaño, and R. C. Anderson, “Onboard Autonomous Rock Shape Analysis for Mars Rovers,” IEEE Aerospace Conference, Big Sky, Montana, March 2002.
- [9] M. Gilmore, R. Castaño, T. Mann, R. C. Anderson, E. Mjolsness, R. Manduchi, and R. S. Saunders, “Strategies for Autonomous Rovers at Mars,” *Journal of Geophysical Research*, vol. 105, no. E12, pp. 29,223–29,237, December 2000.
- [10] V. Gor, R. Castaño, R. Manduchi, R. Anderson, and E. Mjolsness, “Autonomous Rock Detection for Mars Terrain,” Space 2001, American Institute of Aeronautics and Astronautics, Albuquerque, New Mexico, August 2001.
- [11] V. C. Gulick, R. L. Morris, M. A. Ruzon, and T. L. Roush, “Autonomous Image Analysis during the 1999 Marsrokhod Rover Field Test,” *Journal of Geophysical Research*, vol. 106, no. E4, pp. 7745–7764, 2001.
- [12] T. S. Lee, “Image Representation Using 2D Gabor Wavelets,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 18, no. 10, pp. 959–971, 1996.
- [13] E. L. Lehmann, *Nonparametrics: Statistical Methods Based on Ranks*, San Francisco, California: Holden-Day, 1975.
- [14] P. Smyth, “Clustering Using Monte-Carlo Cross-Validation,” *Proceedings of the 2nd International Conference on Knowledge Discovery and Data Mining*, AAAI Press, Portland, Oregon, pp. 126–133, 1996.