Survey on Detection and Segmentation of Rocks on Mars

[1] Mahendra Gaikar, [2]Manthan Jadhav, [3] Cavin D'Souza, [4] Dr. Deepa Parasar
[1,2,3,4] Dept. of Computer Engineering, Saraswati College of Engineering, Kharghar, Mumbai, Maharashtra.
[1]maudeega27@gmail.com [2]manthan.jadhav69@gmail.com [3]cavin786@gmail.com [4]deepaparasar@gmail.com

Abstract: This paper presents a survey on different methods used for detection and segmentation of rocks on mars. Geologists and planetary scientists will benefit from methods for accurate segmentation of rocks in natural scenes. However, rocks are poorly grouped for current visual segmentation techniques- they demonstrate diverse morphologies and have no uniform property to distinguish them from background soil. To deal with this challenge a novel detection and segmentation method incorporating features from multiple scales. These features include local attributes such as texture, object attributes such as shading and two dimensional shape, and scene attributes such as the direction of illumination. The method uses a super pixel segmentation followed by region-merging to search for the most possible groups of super pixels. A learned model of rock appearances identifies whole rocks by scoring candidate super pixel.

1. Introduction

Planetary rovers like those on Mars today can collect more data than can be examined in mission-relevant time frames; they have usually moved on before their images of a site are thoroughly analyzed. The Mars Exploration Rovers Spirit and Opportunity have already amassed a catalog of hundreds of thousands of images. Future rovers will have even greater mobility and lifespan; missions could cover kilometers every day and last for years. This missions crosses vast areas and collect various images of the planetary surface. However, improved exploration speed and efficiency will reduce the time window for any data analysis that is relevant to the ongoing exploration. The identification of observed rocks is an important task in route planning and geologic analysis. Rock shape, weathering, and dispersion carry important information about environmental characteristics and processes. Previous research has produced several automated algorithms for finding rocks, including template-based approaches, stereo geometry, finding closed shapes with an edge detector, and classifying consistent regions with a belief network. These algorithms can detect rocks but rarely find their actual boundary curves.

Segmenting rocks is difficult because they are non-uniform: their texture, color and albedo varies across their surfaces and from one rock to the next. Often Mars images display strong directional lighting with cast shadows and highlights that go against the uniformity assumption. Weak boundary edges must be indirect from context. This work proposes a new segmentation method that utilizes multi-scale signals simultaneously. It aims for a segmentation that is consistent with local features such as texture, object features such as shape and shading, and scene features such as the direction of illumination. The identification of observed rocks is an important task in route planning and geologic analysis. Rock shape, weathering, and distribution carry important information about environmental characteristics and processes.

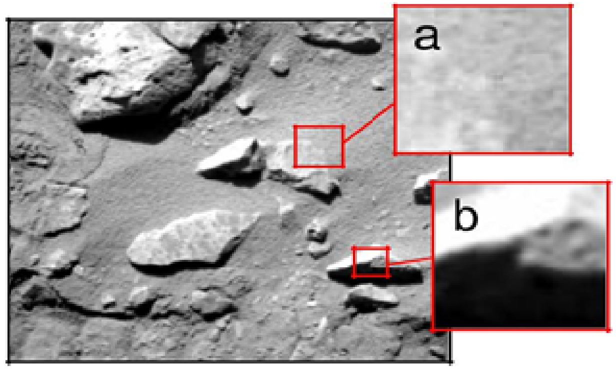

Figure 1 Rocks diverse textures, variable morphologies, and directional illumination all complicate the segmentation problem. Local cues are insufficient: real borders are often weak (a), while the rocks themselves are non-uniform with false interior borders

(b). Image from the Mars Exploration Rover "Spirit," courtesy of NASA/JPL

2. Approach

Rock identification method fragments the source image using a normalized-cuts strategy and merges the resulting fragments, or super pixels into candidate rock regions. A search process identifies high-possibility super pixel groupings where labels are most consistent with the local and global features of our rock model. The search evaluates each possible segmentation of candidate regions using a learned model trained to recognize rocks and soil under directional lighting.

2.1 Creating Super pixels

An initial over-segmentation fragments the image into a set of super pixels — areas of uniform properties that become the atomic elements of the search. The use of super pixels instead of image pixels reduces the space of possible segmentations while still ensuring that some valid segmentation captures the shapes of all rocks.

2.2 Modeling Rocks

The method gathers rocks from one or more super pixels that form a closest region in the image. In order to identify quality segmentations, we learn a mapping from an image region's features to the possibility that the region is actually a rock. Many of the features we consider for this application are similar to those used in region merging by Mori[5] . A brief description of the feature set follows. Rocks have very unequal shapes; no two exhibit exactly the same morphology. However, their shapes are recognizably different from a randomly generated region. They tend to be more ellipsoidal and convex than random groupings of super pixels. Dunlop [1] provides a widespread synopsis of these shape measures and their utility in describing rocks.

2.2.1 Shape

Rocks have very irregular shapes; no two display exactly the same morphology. However, their shapes are recognizably different from a randomly generated region. They tend to be more ellipsoidal and convex than random groupings of super pixels. Dunlop provides a widespread summary of these shape measures and their utility in describing rocks[1].

2.2.2 Texture

Our texture measure uses the method of Varma and Zisser-man[18]. We combine all images in a training set with the Maximum Response 8 filter bank. This results in an 8-dimensional response vector for each pixel. The responses are totaled for several training images and collect them using k-means to form a set of 16 universal textons. We use these textons to compute a texton map for each novel image by allocating pixels to their Euclidean-nearest texton.

2.2.3 Shading

The used shading calculation develops several unique features of the Mars problem domain. First, a precise approximate of sun direction is available for each image in the catalog — the illumination direction can be observed directly using an on-board sensor or calculated from rambler pose and ephemeris data. Additionally, rocks on the Martian surface are generally Lambertian. This makes them agreeable to the classic shape-from-shading reflectance formula. However, shape-from-shading techniques generally perform poorly on non-synthetic data .

2.3. Searching Segmentations

Each candidate segmentation defines a set of rock regions in the image. For our purposes, every pixel not belonging to some rock is considered to be soil. Thus, each segmentation determines a set of rock regions $X = \{x_1, x_2 . . ., x_n \}$ and soil pixels $Y = \{y_1, y_2, . . . , y_n \}$. A segmentation's score depends on how well these match the discovered appearances of actual rocks and soil. Class possibilities is generated using two Support Vector Machine (SVM) classifiers with radial basis kernel functions [6]. The rock SVM uses the complete 26-dimensional feature vector describing both local and global properties of the region. It learns the mapping from this feature vector to the probability $P_r (x_i)$ that a region is a rock. The soil SVM produces an estimate $P_s (y_i)$ that each pixel is soil. It generates these possibilities independently for each super pixel using local texture histograms alone. Averaging soil possibilities across different scales of super pixels yields a value $P_n (y_i)$ for each image pixel.

In addition to its role in scoring candidate segmentations, the soil classifier acts as a heuristic for initializing the search procedure. The space of possible rock regions is too large for an exhaustive search. Under these conditions, data-driven heuristics that favor more probable super pixel groupings can dramatically improve performance. We group super pixels with a low probability of being soil into contiguous regions that become rocks for the first round of search. This initial guess based on texture dramatically improves performance.

The search procedure seeks higher-scoring segmentations by modifying rock regions in one of four ways: (1) Growing the rock by adding a super pixel at some scale; (2) Shrinking the rock by subtracting a super pixel at some scale; (3) Splitting the rock by removing a super pixel to make a new rock region; and (4) Merging two contiguous rock regions into a larger rock region. This search procedure generates regions that could not have been constructed using super pixels from a single scale exclusively.

The search seeks a segmentation that maximizes the following objective function:

$$f(X,Y) = A(x_i)P_r(x_i) + P_n(y_i)$$

where A(x) is the pixel area of a rock region. Rock and soil areas are zero-sum so the two models compete to explain the scene.

3. Methodologies for Detection and Segmentation

Analysis of multi-scale segmentation method using images of the Martian planetary surface is presented here. It determines the detection accuracy, compare the value of individual features, and measures the error in boundary localization.

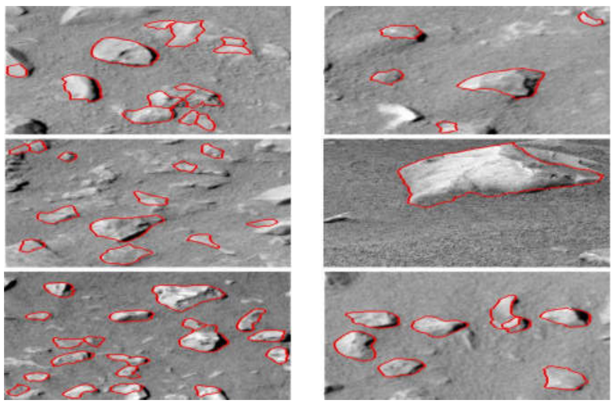

Figure 2: Detected rocks

The involved images drawn from the Spirit Mars Rover Panoramic Camera image catalog at various locations along its traverse path in Gusev Crater. This study is restricted to near-field images of the terrain directly in front of the rover. These downward-looking images were favored for two reasons. First, for uniform distributions of rocks, near-field scenes appeared less cluttered due to a lack of foreshortening. Second, near-field rocks contained a larger number of pixels for computing texture and shading statistics. Creating multi-scale super pixel segmentations for each image was the most computationally-intensive portion of the algorithm, taking several days to compute for the whole data set. Figure above shows some examples of rocks detected. The following sections provide a quantitative performance evaluation of presented method.

3.1 Performance Results

A variety of system performance measures investigated here as, First, the influence of individual features on region classification accuracy.

3.1.1 Feature Comparison

The Venn diagram of Figure below shows the influence of shape and shading features on rock-labeling accuracy. The SVM has been trained for each set of training features and evaluated its accuracy in predicting rock versus non-rock regions. A framework search over training parameter values identified finest training parameters for each feature subset.

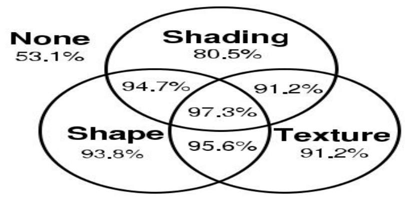

Figure3: Labeling accuracy for different attribute sets. The *None* feature corresponds to a blind policy that always chooses the most frequent class.

3.1.2 Rock Detection Accuracy

Several different performance measures for rock detection and segmentation [8, 4] are there. "Area of overlap criterion" to reject detections whose borders do not accurately match those of the true rock. Match detection is done to their most similar true rocks; successful detections are those whose intersection area comprises 50% or more of their union. Based on Precision, recall scores is generated by varying the confidence entrance for preserving rocks in the final segmentation; this value can be tuned to persuade high-precision or high-recall behavior. The *precision* is the percentage of true rocks among candidate segmented rocks, while the recall is the percentage of rocks segmented from the total actually present in the scene. The precision and recall scores for a test set of 56 images containing over 230 rocks. There are relatively few rocks per image, so performance varies significantly from one image to the next. The standard deviation for search performance generally ranges between 0.2 and 0.3 for both recall and precision. Nonetheless, the test suggests that the segmentation search provides a reserved improvement in both precision and recall. This confirms visual evidence that the configuration search improves the reliability of the segmentation.

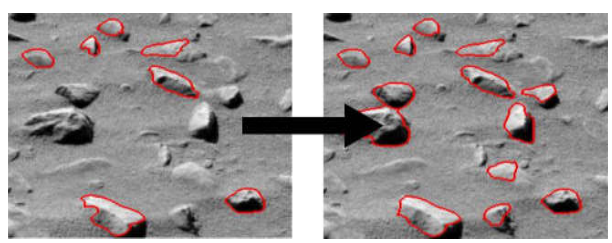

Figure 4: The effect of the configuration search on rock segmentation results. The initial segmentation (left) loses rocks whose probabilities fall below the confidence threshold. The search process recovers their boundaries and improves their score (right).

3.1.3 Boundary Localization Accuracy

The Chamfer distance measures how closely the detected rock borders match the actual rock borders. This value is the average distance from a detected region boundary pixel to the closest boundary pixel on the real rock region. Only shape matches from successful detections are shown. For studied data set the search procedure provides no apparent advantage for boundary localization beyond that which is required to detect the rocks in the first place. In other words, the search increases the number of rocks that are discovered but it does not improve their boundary localization accuracy.

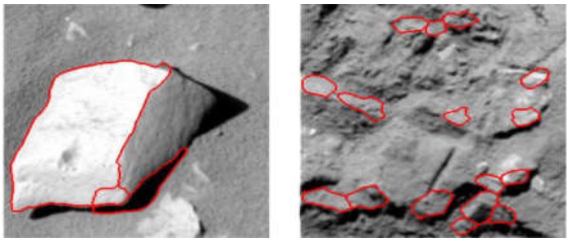

Figure.5: Two failure modes: large rocks (left) and structured soil (right).

4. Conclusion

Several scene factors manipulated the quality of resulting detections. Figure above shows some examples of failure modes. Large rocks are rare in the training set; the method often divides them into pieces or fails to find them altogether. The greedy search cannot find the largest rocks; these re-quire many super pixel additions. Such piecemeal construction involves short-term sacrifice - a non-ellipsoidal shape, for example - that the hill-climbing approach avoids. Structured soil also proved difficult to disambiguate. The rightmost image shows soil that has been compressed by the rover's wheel; there are no rocks in the image but the system returns several false positives due to bogus shadows and patches of odd texture.

Despite these barriers, our multi-scale segmentation technique provides promising results for object identification in natural scenes. This work draws together a wide range of recent advances in object recognition, including super pixel representations, texton analysis, and simultaneous segmentation and recognition. The work demonstrate the applicability of these ideas to the difficult problem of rock detection and segmentation. Rocks in Mars images have the unique challenges of directional lighting, variable object morphology and weak local prompts. The work shows how such objects can be segmented by combining local, object-level and scene-level features in a multi-scale technique. . The super pixel segmentation used here was selected after a qualitative comparison with others readily available; a more thorough analysis of different segmentation methods may also be appropriate.

References

[1] H. Dunlop, Automatic Rock Detection and Classification in Natural Scenes, vol. CMU-RI-TR-06-40. August 2006.

[2]. V. Gor, R. Castano, R. Manduchi, R. C. Anderson, and Mjolsness, "Autonomous rock detection for mars terrain," in Proceedings of AIAA Space 2001, August 2000.

[3]. A. Castano, R. C. Anderson, R. Castano, T. Estlin, , and Judd, "Intensity-based rock detection for acquiring on-board rover science," in Lunar and Planetary Science, no. 35, 2004.

[4] D. R. Thompson, S. Niekum, T. Smith, and D. Wettergreen, "Automatic detection and classification of geological features of interest," in IEEE Aerospace Conference Proceedings, Big Sky Montana, March 2005.

[5] G. Mori, X. Ren, A. Efros, and J. Malik, "Recovering human body configurations: Combining segmentation and recognition," in IEEE International Conference on Computer Vision, 2004.

[6] C.-C. Chang and C.-J. Lin, LIBSVM: A Library for Support Vector Machines, 2001.

[7] M. Varma and A. Zisserman, "A statistical approach to tex-ture classification from single images," in International Jour-nal of Computer Vision: Special Issue on Texture Analysis and Synthesis, vol. 62:1-2, pp. 61–81, 2005.

[8] D. R. Thompson and R. Castano,˜ "Automatic detection and classification of geological features of interest," in IEEE Aerospace Conference Proceedings, Big Sky Montana , March 2007.