

Multi-scale Features for Detection and Segmentation of Rocks in Mars Images

Heather Dunlop, David R. Thompson and David Wettergreen, The Robotics Institute, Carnegie Mellon University


1. Introduction

Automatic detection and segmentation of rocks will be increasingly important in future missions to Mars and beyond. Future rovers could travel kilometers per day and last for years. Even today the Mars Exploration Rover catalog includes hundreds of thousands of images.

Rock shape, weathering, dispersion carry important information about environmental characteristics and processes. Automated methods can allow automatic extraction of geologic information by enabling the real-time calculation of rock size, shape and class distributions.

Challenges:

- Segmenting rocks is difficult due to non-uniform texture, color and albedo.
- An accurate rock boundary is essential for further geologic analysis.
- Mars exhibits strong directional lighting with cast shadows and highlights.



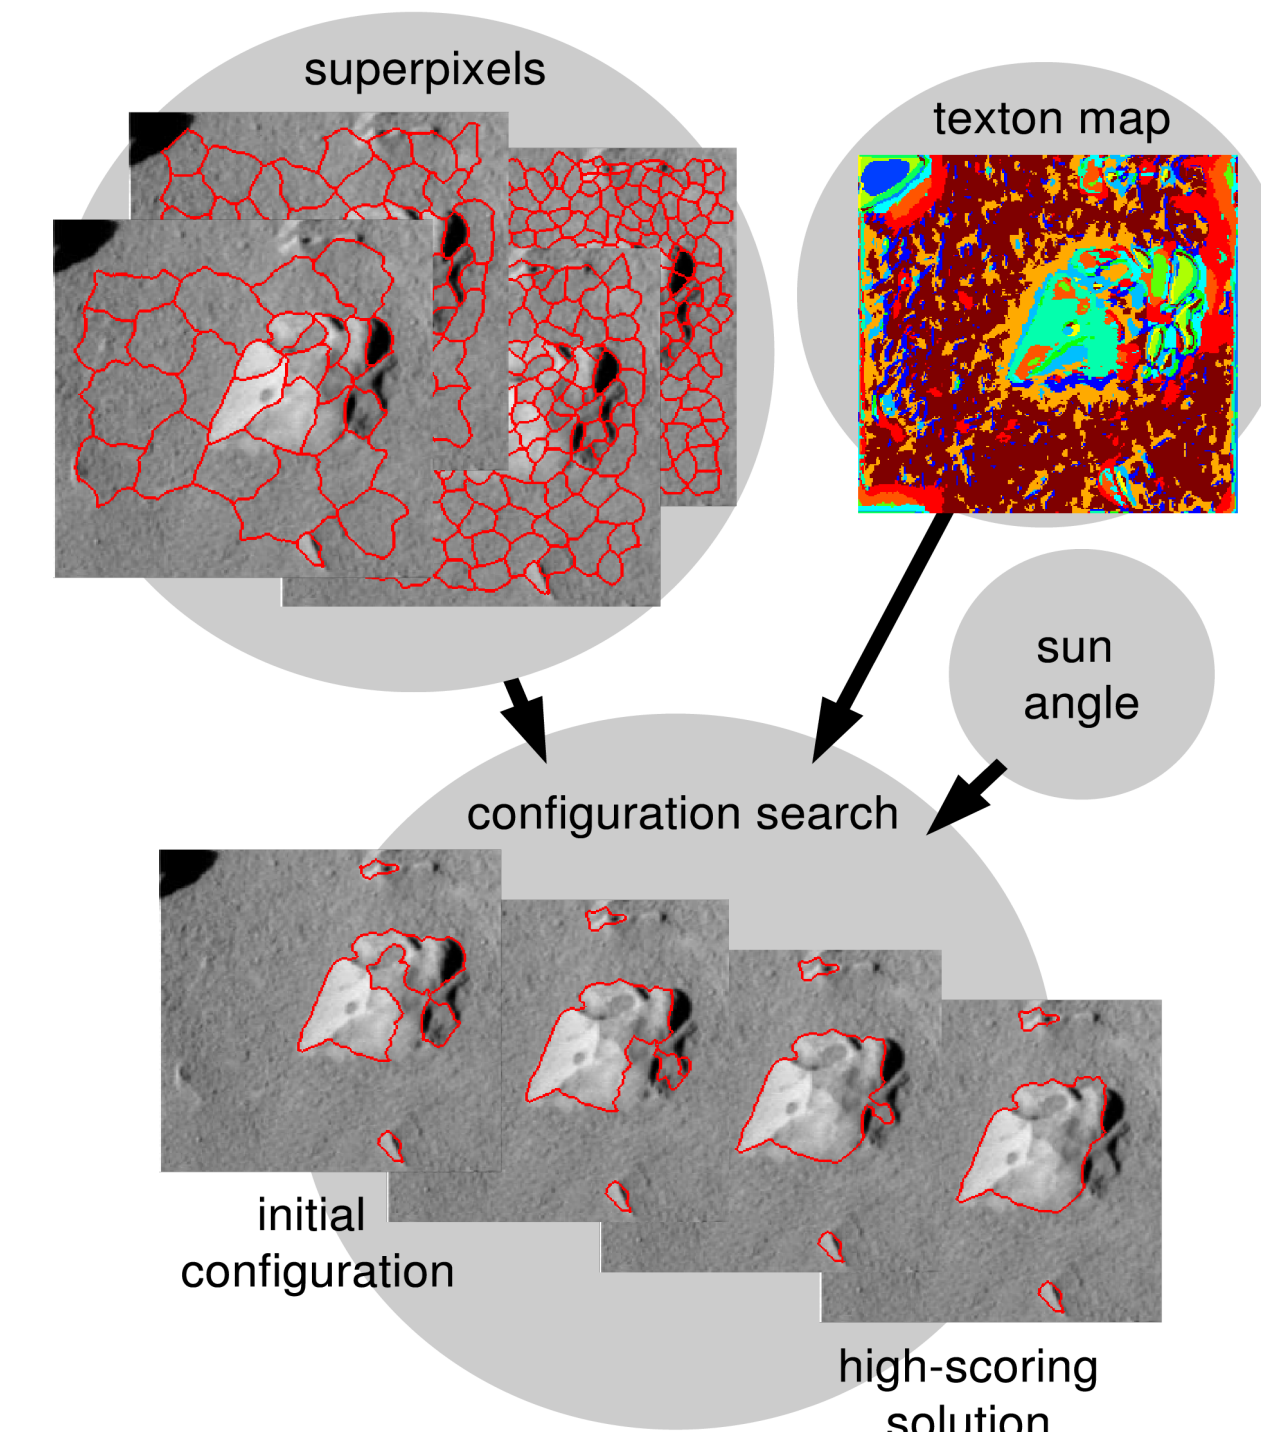
Rock segmentation difficulties: a) rock borders are often weak and b) rocks themselves are non-uniform with false interior borders.

2. Approach

1) Fragment the image using normalized-cuts at four different scales.

2) Merge the resulting superpixels into candidate regions [1].

- A learned **rock model** recognizes rocks and soil under directional lighting.
- A **greedy search** finds superpixel groupings that are most consistent with local and global features of the rock model.



3. The Rock Model

A Support Vector Machine (SVM) model maps a group of superpixels onto the probability that the region is a rock. We use the following features:

Shape Features

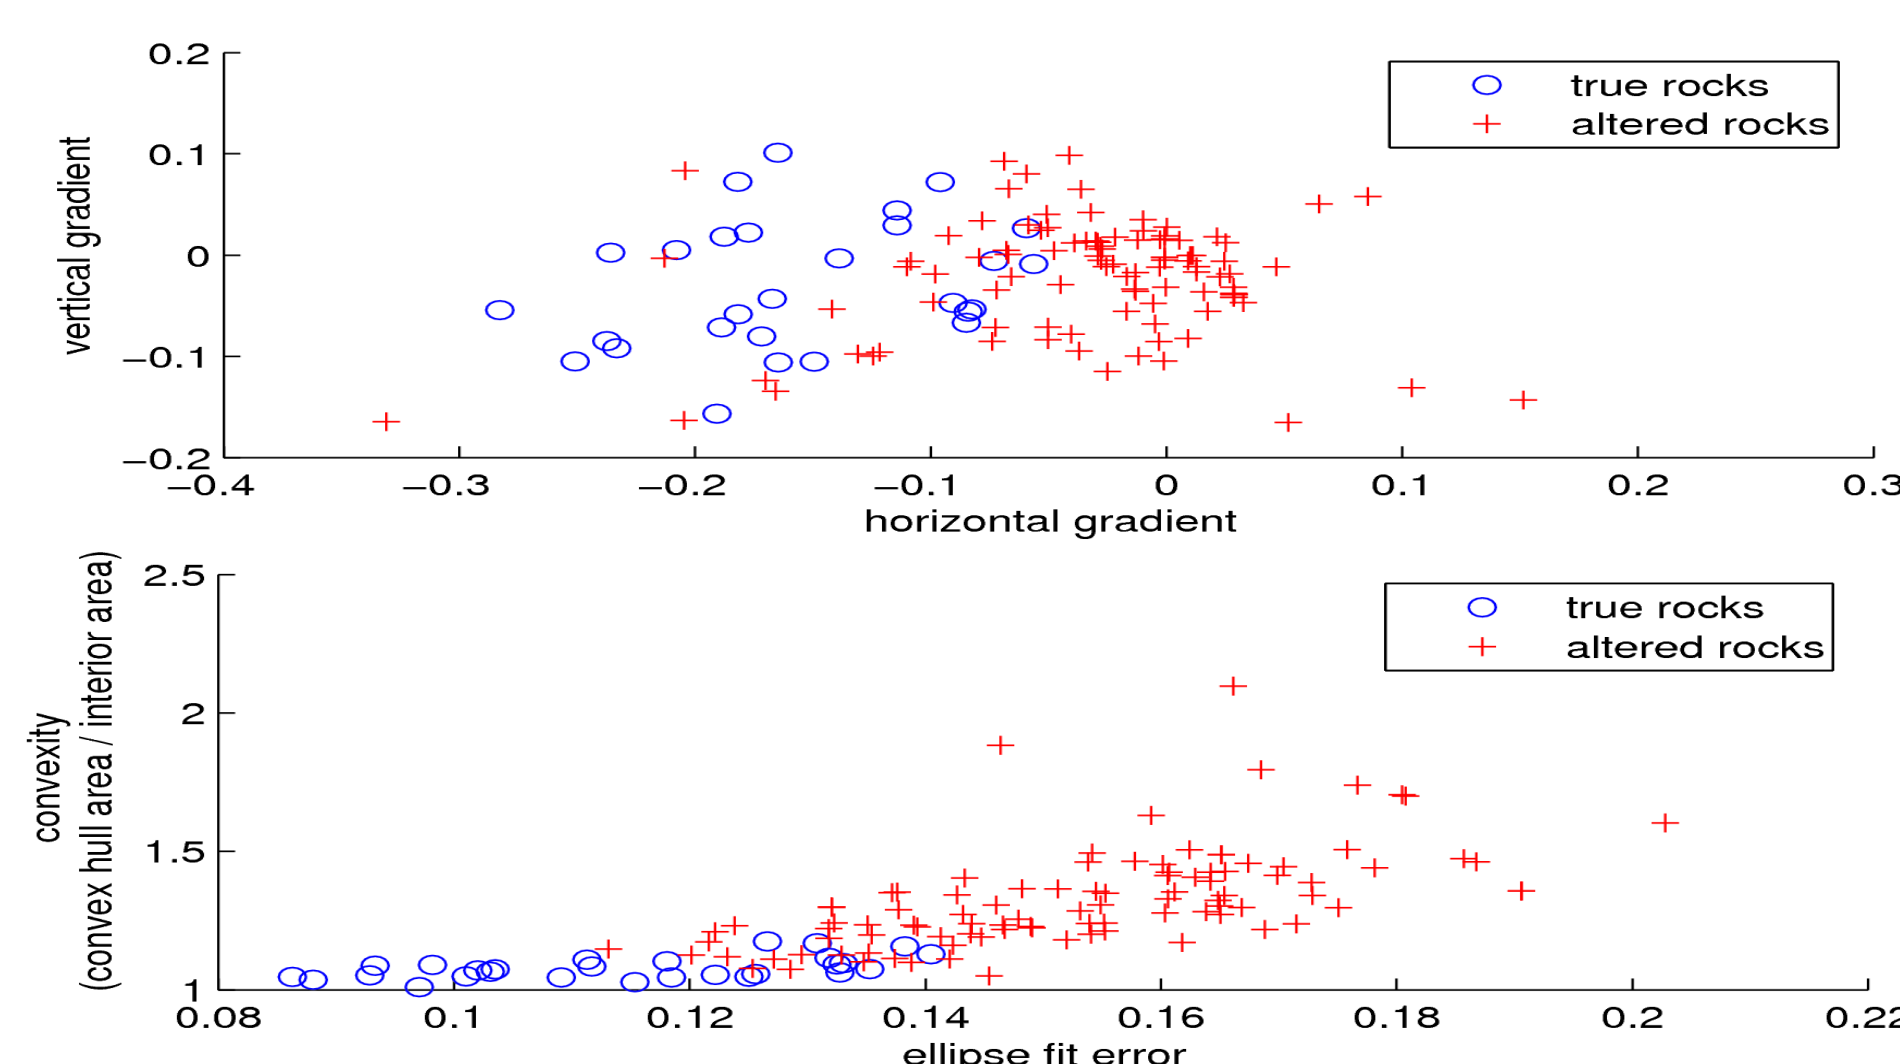
- *Ellipse error*: residual error between the rock region's contour and its corresponding best-fit ellipse.
- *Convexity*: ratio of the convex hull perimeter (area) to the actual perimeter (area) of the rock.
- *Circularity*: square of region's perimeter divided by its area.

Texture Features

- Histogram of universal texton counts from the region [2].
- Chi-squared distance between texton histogram of interior region and that of its local neighborhood.

Shading

- Azimuth and elevation of the sun in the camera frame.
- 2-D intensity gradient on the rock's surface.



4. Searching Segmentations


Rock and Soil Models

- Each candidate segmentation determines a set of rock regions $X = \{x_1, x_2, \dots, x_n\}$ and soil pixels $Y = \{y_1, y_2, \dots, y_n\}$. Segmentations' scores reflect how X and Y match the learned appearance of rocks and soil.
- A "Rock SVM" maps shape, texture and shading feature vectors to the probability $P_r(x_i)$ that a region is a rock.
- A "Soil SVM" learns the mapping from texture feature vector to the probability $P_s(y_i)$ that each pixel is soil.


Greedy Search

- Initial configuration: group superpixel with low probability of being soil into contiguous regions that become rocks.
- Seeks higher scoring segmentations by modifying rock regions by growing, shrinking, splitting or merging superpixels.
- Maximizes the objective function:

$$f(X, Y) = \sum_{x_i \in X} A(x_i) P_r(x_i) + \sum_{y_i \in Y} P_s(y_i)$$



Typical segmentation results using local texture features (left) and the entire feature set (right). Global features like shape provide important cues to localize rock boundaries.



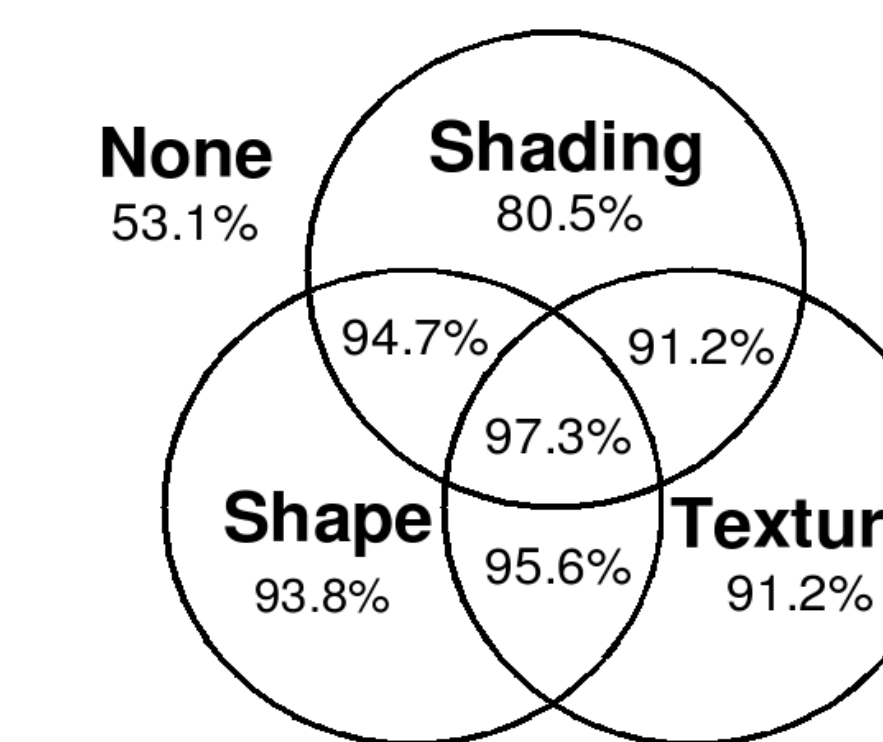
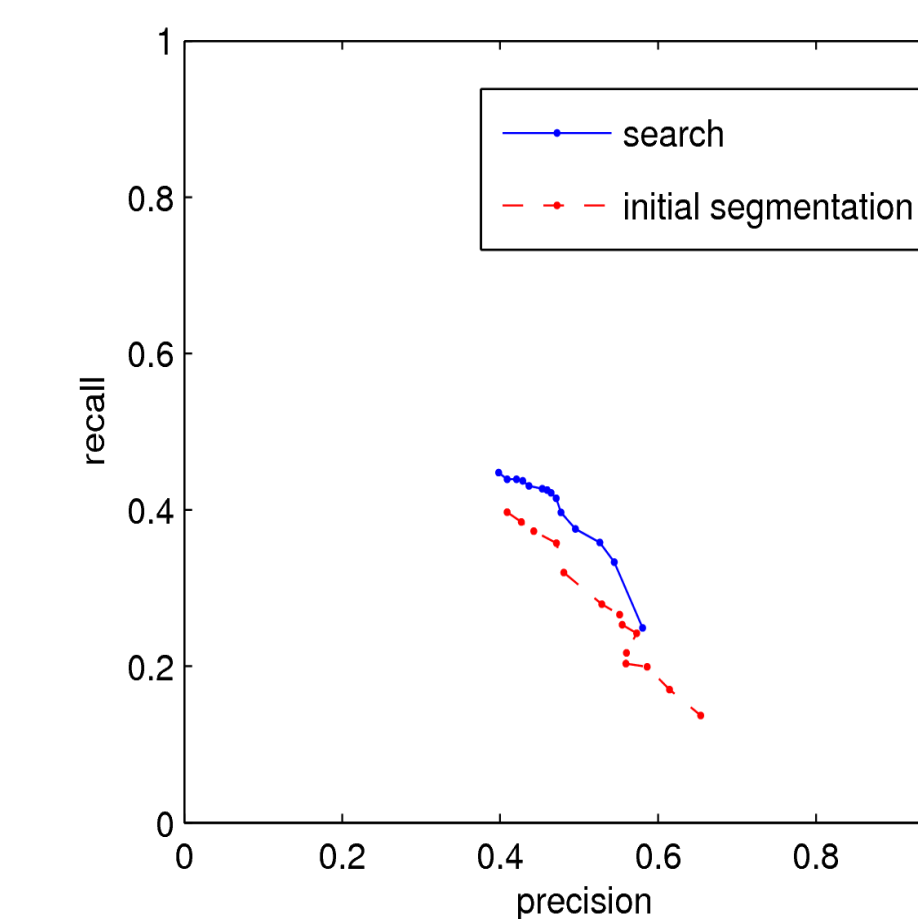
The effect of the configuration search: the initial segmentation (left) loses rocks whose probabilities fall below the confidence threshold; the search process (right) recovers their boundaries and improves their scores.

5. Experiments


Our test set used Images drawn from the Spirit Mars Rover Panoramic Camera image catalog at various locations along its traverse path in Gusev Crater. It included 150 hand-labeled images, each 256 x 256 pixels, with over 500 Mars rocks and diverse lighting.

Evaluation method

- Match detections to their most similar true rocks; successful detections are those whose intersection area constitutes 50% or more of their union.
- Generate precision and recall scores by varying the confidence threshold for preserving rocks in the final segmentation.
- Compare results with and without configuration search.



Region labeling accuracy with various feature sets.



References: [1] X. Ren and J. Malik, "Learning a classification model for segmentation," *ICCV* 2003. [2] M. Varma and A. Zisserman, "A statistical approach to texture classification from single images," *IJCV* 62:1-2, 2005.

This research was supported by NASA under grants NNG0-4GB66G and NAG5-12890.