

# Automatic Rock Detection and Classification in Natural Scenes

Heather Dunlop

CMU-RI-TR-06-40

August 2006

Robotics Institute  
Carnegie Mellon University  
Pittsburgh, Pennsylvania 15213

© Carnegie Mellon University

*Submitted in partial fulfillment of the requirements  
for the degree of Master of Science*



## Abstract

Autonomous geologic analysis of natural terrain is an important technique for science rovers exploring remote environments such as Mars. By automatically detecting and classifying rocks, rovers can efficiently survey an area, summarize results and concentrate on unusual discoveries, thus maximizing the scientific return. Similar techniques can be used in natural settings on land and under water where characterization of context is important to robot behavior.

In this work, techniques for characterizing rocks using albedo, color, texture and shape properties are developed. These and other features are used to detect rocks, using a superpixel segmentation algorithm to locate and delineate the rocks with an accurate boundary. Once segmented, a Bayesian procedure is applied to geologically classify rocks. This capability is fundamental in enabling robots to conduct scientific investigations autonomously.

Experiments demonstrate the usefulness of the albedo, color, texture and shape measures and establish the accuracy in computing two important geologic shape metrics, sphericity and roundness. The high performance of the rock detection and segmentation algorithm is analyzed and the results of geologic rock classification are presented.

These methods successfully perform geologic rock analysis, while giving an indication of what features are most important for this and similar tasks. As color, texture and shape are such powerful features, this approach is also applicable to detecting, segmenting and classifying other natural objects and terrain features.

Thesis supervisor: Prof. David Wettergreen



## Acknowledgments

This work was made possible by my advisor, David Wettergreen. Thank you for your guidance throughout this project, and for introducing me to a fascinating application of computer vision to space robotics. I also want to extend my gratitude to my other committee members, Alyosha Efros and David Thompson. The rock detection and segmentation work presented in this report began as a course project with David Thompson, guided by Alyosha as our course professor.

Thank you to Geb Thomas and Ingrid Ukstins Peate at the University of Iowa for providing the data sets on which all experiments were performed.

My fellow RoboGrads are a great bunch and I want to thank them for making the last year in Pittsburgh so enjoyable. I am also grateful to my parents for doing everything that parents do. But most of all, thank you to Rob Couture for introducing me to the field of computer vision, and for always providing the inspiration and encouragement that I needed to get this done.



# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Motivation . . . . .	1
1.2	Scope . . . . .	2
1.3	Approach . . . . .	2
1.4	Overview . . . . .	3
<b>2</b>	<b>Feature Extraction</b>	<b>4</b>
2.1	Albedo and Color . . . . .	4
2.1.1	Related Work . . . . .	4
2.1.2	Methods . . . . .	5
2.2	Texture . . . . .	7
2.2.1	Related Work . . . . .	8
2.2.2	Methods . . . . .	9
2.3	Shape . . . . .	14
2.3.1	Related Work . . . . .	15
2.3.2	Methods . . . . .	16
<b>3</b>	<b>Rock Detection and Segmentation</b>	<b>24</b>
3.1	Related Work . . . . .	27
3.2	Approach . . . . .	27
3.2.1	Superpixel Segmentation . . . . .	28
3.2.2	Feature Extraction . . . . .	34
3.2.3	Rock Detection . . . . .	37
<b>4</b>	<b>Rock Classification</b>	<b>39</b>
4.1	Related work . . . . .	39
4.2	Approach . . . . .	41
<b>5</b>	<b>Experiments</b>	<b>42</b>
5.1	Data Set . . . . .	42
5.2	Rock Features . . . . .	42
5.2.1	Albedo and Color . . . . .	42
5.2.2	Texture . . . . .	43
5.2.3	Shape . . . . .	45
5.3	Rock Detection and Segmentation . . . . .	53
5.3.1	Region Labeling Accuracy . . . . .	53
5.3.2	Rock Detection Accuracy . . . . .	54
5.3.3	Boundary Localization Accuracy . . . . .	57
5.4	Rock Classification . . . . .	59
<b>6</b>	<b>Conclusions</b>	<b>59</b>
6.1	Summary . . . . .	61
6.2	Contributions . . . . .	61
6.3	Future Work . . . . .	62
6.4	Closing Remarks . . . . .	63



# 1 Introduction

This work addresses the problem of automatically analyzing images of rocks and soils, as would be seen by a planetary rover. The goal is to develop algorithms and methods to accurately locate rocks in a natural scene and to classify the rocks into geologic types. Locating objects involves detecting them and then separating, or segmenting, them from the background. This is a difficult task because of the loose definition of a rock and the complexity of the natural scene. Rocks come in all colors, shapes and sizes, with a variety of distinguishing features. An example scene is shown in Figure 1. Finding appropriate methods for rock detection and classification requires the comprehensive examination and experimentation of computer vision techniques undertaken in this work.

## 1.1 Motivation

Distinguishing types of rocks in natural scenes, such as the desert terrain of Mars, presents many challenges. Yet if rocks can be located and their geologic characteristics determined, it will greatly improve the efficiency of future rover explorers in remote environments.

Currently, images and data recorded by rovers on Mars are transmitted back to Earth where scientists perform analysis and upload a new set of instructions for the rover to follow. However, the bandwidth of data transmitted from the rover back to Earth is limited. Scientists spend large amounts of time examining options and making choices for instrument observations based on the limited time and bandwidth available. They cannot collect all the data that they desire.

If the rover could make intelligent decisions about what observations to make, how to summarize



Figure 1: An example natural rock scene.

them, and what to send back to Earth, it would greatly improve efficiency and increase the scientific return. For example, if a geologist is interested in a specific type of rock in a given area, the rover could seek that type, summarize its observations and return the best data related to that type of rock.

In addition, the rover returning data to Earth and then waiting for a new set of commands to follow is an inefficient process. With autonomous geologic analysis of a region, a rover could make intelligent decisions without waiting for specific commands from scientists on Earth.

Rock detection is also important for obstacle avoidance and path planning. Locating rocks, and perhaps estimating terrain properties from a geologic perspective, could better inform rover navigation algorithms.

Each of these applications can improve the operations of planetary rovers or other remote exploration vehicles.

## 1.2 Scope

This thesis tackles the problem of locating and classifying rocks. Accurate boundary localization is required, as is a high precision and recall of detected rocks. Rock attributes such as albedo, color, texture and shape must be characterized in a manner that provides a useful analysis of these properties to geologists. A full classification system will take the computed rock features and accurately classify rocks geologically.

As this technology is still new, the goal here is to develop algorithms that effectively perform the tasks of rock detection and classification, with efficiency being secondary. As such, the current system is more applicable to the post-processing of images from Mars rovers, rather than on the rovers themselves. This system is still of great use because of the large set of images already returned from the Mars Exploration Rovers (more than 160,000 to date). It is not feasible for geologists to analyze each image, but a computer can take the time to perform this task, improving the geologic analysis of regions already explored. Future systems using similar techniques will be able to perform this analysis on Martian rovers.

## 1.3 Approach

Detecting rocks in an image is the first step in the process. A rock is not a clearly defined object. Anything from a tiny pebble to a large boulder is a rock. Rocks may also be covered partially by sand or occluded by other rocks. The composition of the rock may even be the same as the substrate. Variations in illumination and shadows also present challenges. Standard object recognition approaches may not work because of the variety of colors, textures and shapes of rocks. Localizing the rock boundary with pixel accuracy may be necessary for further analysis of rocks. For this, high performance segmentation algorithms are necessary to accurately partition the image.

The three main geologic classes of rocks are igneous, metamorphic and sedimentary. Within each of these, geologists further distinguish different rock types. Color, texture and shape properties can be used to characterize different classes.

Color and albedo properties can help in identifying carbonaceous materials or oxidized iron in rocks. They can be represented by computing the mean and variance or in the form of histograms of the distribution of color or intensity over the rock.

For geologists, important properties of rock texture are granularity, directionality, polish and markings such as striations or percussion marks. Texture can be represented in different ways, some derived from older approaches in geology and some from newer techniques in computer vision. The fractal dimension and statistics of the gray level co-occurrence matrix are older methods, while newer approaches tend to involve convolving the image with a set of filter banks.

Rock shape can be categorized into form and roundness properties. Form represents the overall shape of the rock, while roundness measures the sharpness of its corners. Rock shape can pro-

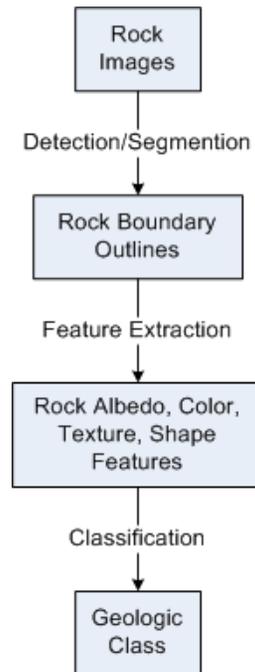


Figure 2: Overview of the rock detection, segmentation and classification process.

vide geologists with information about a rock’s abrasion history, including the method and distance of transport. Traditionally, sphericity and roundness are used to characterize rock shape. However, other shape metrics can also provide useful information, perhaps by characterizing form and roundness in a more accurate or more informative manner. The accuracy of automated methods of computing sphericity and roundness are important.

With the combination of rock color, texture and shape features, we can classify rocks geologically using standard techniques in machine learning. This is done with the use of a set of rocks classified by geologists. These geologic classes are also used to assess the effectiveness of each of the color, texture and shape measures.

The final system achieved here has the ability to detect, segment and geologically classify rocks. This process is shown in Figure 2. Rock boundary outlines are first detected using a two-step segmentation and then detection process. Albedo, color, texture and shape features are evaluated on each rock. Finally, a classifier is applied to predict the geologic class of rocks.

## 1.4 Overview

This report explores techniques for characterizing, detecting and classifying rocks. Section 2 describes methods for extracting features to characterize rock albedo, color, texture and shape. These features will be used to detect rocks as well as to classify them. Section 3 details the algorithm for segmenting images and detecting rocks. In Section 4 we explore methods for performing geologic classification on rocks. Experiments on all rock detection and classification methods are described and results given in Section 5. Studies are performed on each rock property individually, as well as their use in the detection and classification process. Section 6 provides a summary and outlines some conclusions and future work.

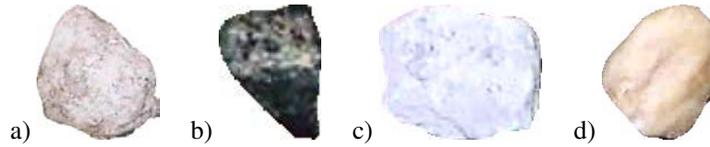


Figure 3: Sample rocks showing the variations in coloring: a) brown limestone, b) dark volcanic igneous, c) light limestone with calcite crystals, d) reddish limestone.

## 2 Feature Extraction

Fundamental properties of rocks are size, shape, texture and mineralogical composition [31]. Albedo and color properties can provide useful information about the composition of rocks. Rock surface texture gives clues to its crystal content and abrasion history. Shape provides information about the origin and transport of rocks. These properties can all be characterized from images of rocks. This section looks at how these attributes are used in geology and details methods for their automated analysis from images. Some of these methods have been applied to rock geology before, while others are newer algorithms developed in the field of computer vision.

In this section it is assumed that we have a boundary for a region which we are analyzing. In Section 3 it will be shown how these albedo, color, texture and shape features can help identify image regions containing rocks and produce an accurate segmentation of the rock boundary. In Section 4 we will use the rocks detected and segmented and the rock features described in this section to geologically classify rocks. Thus, these rock properties are computed from the original image and a binary mask showing the location of the rock or potential rock.

### 2.1 Albedo and Color

In the field, geologists specify color by reference to a chart published by the Geologic Society of America or by the Munsell Color Company [2]. The most common colors of rocks are red, brown and yellow which are typically due to the presence of ferric oxide cement, gray-black which reflects the presence of carbonaceous material, and colorless, such as quartz, which contains neither ferric oxide nor free carbon [2]. Color is an especially important characteristic for shale. Black shales have a higher content of organic matter. Red coloring does not necessarily reflect the content of iron, but does indicate its state of oxidation [44]. Local weathering conditions can also have an affect on rock color and can indicate the amount of dust cover. Examples rocks are shown in Figure 3 to illustrate variations in coloring.

In this section we develop methods for the analysis of the albedo and color of rocks. Albedo is defined as the reflectivity of the rock - that is, the ratio of the intensity of the reflected light to the incident light.

There is a limited number of ways in which rock albedo and color can be characterized. Here we describe the use of the statistical measures mean and variance, as well as histograms. These methods are applied to albedo and color.

#### 2.1.1 Related Work

Approaches taken for characterizing rock albedo from images usually start simple, for instance, the average grayscale value of the rock [5, 50]. The variance of the rock intensity can also provide a feature [50]. The hue, saturation and intensity can be divided into bins and form a 2-dimensional histogram [51]. Color can be handled by performing texture analysis independently on hue, saturation and intensity or red, green and blue channels [35]. For a more complete treatment of color,

analysis of the reflectance spectra covering many wavelengths is necessary. This technique with a neural network allowed for the development of a detector for carbonate and other minerals [23].

### 2.1.2 Methods

An image can be represented by the equation

$$I(x, y) = L(x, y)R(x, y)$$

where  $I(x, y)$  is the input image,  $L(x, y)$  is the illumination image and  $R(x, y)$  is the reflectance image for each pixel  $(x, y)$  [60]. We are attempting to describe  $R(x, y)$  when given only  $I(x, y)$ . This is an ill-posed problem as the number of unknowns is twice the number of knowns. We are attempting to characterize the reflectivity of the rock, when all we are given is an image of it. Numerous illumination complications are present here, including highlights on the rock, shadows on the rock, cast shadows on the ground beside the rock or even possibly cast on a rock by another rock or other object. We cannot determine the true reflectivity of the rock. Other complications not even accounted for in this equation include light scattering in the atmosphere and dust or other materials present on the rock.

There is no simple method for tackling this problem. One solution is to place patches of known color in the scene and color-calibrate according to these. However, the best we can do for now is to normalize each intensity image to zero mean and unit variance for pixel intensities to help account for changes in camera parameters when the image was taken. In scenes with highly directional lighting causing highlights and shadows, the computed intensity variance for a rock will be larger than the true value. This cannot be corrected for at this time. For color images, color balancing can be done using Matlab's adaptive histogram equalization routine. This enhances the contrast of the image by operating on small tiles.

#### 2.1.2.1 Intensity Mean and Variance

The simplest method of characterizing rock albedo involves two statistical measures: mean and variance. The mean pixel intensity represents the reflectivity of the rock, while the variance in intensity provides a measure of how uniform the rock reflectivity is.

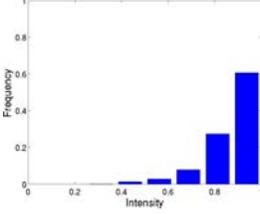
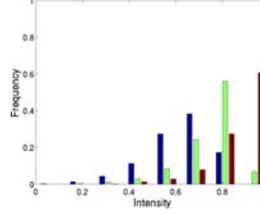
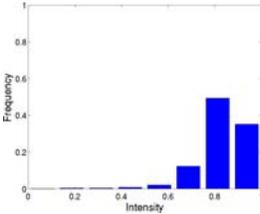
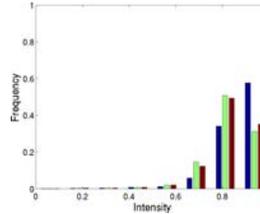
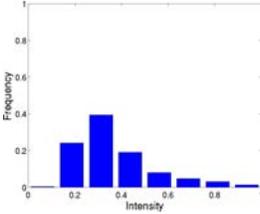
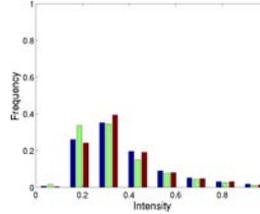
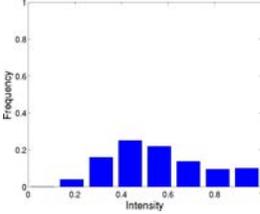
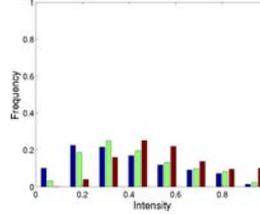
#### 2.1.2.2 Intensity Histogram

A histogram of intensities gives a more complete representation of the reflectivity of a rock. Often a rock will consist of regions of different reflectivity, which intensity mean and variance alone cannot accurately characterize. Rock pixel intensity ranges from zero to one. We divide this range into eight bins to compute the histogram and then normalize so that the elements of the resulting vector sum to one.

#### 2.1.2.3 Color

A complete representation of the color of a rock would involve determining the reflectivity at all wavelengths. However, here we are only dealing with images of a rock which provide the intensity for each pixel at three different wavelengths. Pixel color is commonly represented in one of three ways: Red Green Blue (RGB), Hue Saturation Value (HSV) or CIE  $L^*a^*b^*$  (CIELAB). In RGB, color is represented by three co-ordinates for red, green and blue, each between zero and one. HSV is a non-linear deformation of RGB and can be viewed as a cone with black at the apex. The angular parameter is hue, the distance from the central axis is saturation and the distance from the apex of the cone is value. However, neither the RGB or HSV color space is uniform in that the numerical distance between colors does not correspond to the distance perceived by humans. In particular, we desire a color space in which a change of a certain amount in a color value should produce about the

Table 1: Example rocks and their albedo and color measures.

Rock	Mean Intensity	Variance of Intensity	Intensity Histogram	RGB Color Histogram
	0.86	0.010		
	0.83	0.011		
	0.37	0.027		
	0.56	0.042		

same change in the visual importance. CIELAB is such a color space.  $L^*$  represents lightness,  $a^*$  gives the position between red and green, and  $b^*$  gives the position between yellow and blue.

To represent rock color, RGB, HSV or CIELAB may be used. To characterize rock color, we can then compute the mean and variance over each color channel, as was previously done for intensity. For a more complete color representation, a color histogram can be used. The previously explained method of intensity histograms can be applied to each color channel, resulting in a 2-D histogram.

#### 2.1.2.4 Summary

Each of these methods are able to characterize the albedo and color properties of rocks. Examples are shown in Table 1 which illustrate the use of each method.

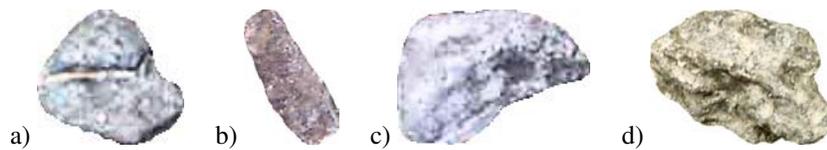


Figure 4: Sample rocks showing the variations in texture: a) two amphibolite pebbles cemented together (cementation line visible), b) sandstone with fine-grained quartz, c) and d) limestone with abrasion marks visible.

## 2.2 Texture

Surface texture is the size, shape and arrangement of the component elements of a rock, as well as surface markings such as polish, striations and pits. Properties of crystals in the rock, such as grain-size, distribution, sorting, permeability, shape and orientation are also important [2].

There are two types of sedimentary rocks. Clastic rocks have grains that were mechanically deposited while nonclastic rocks have grains that were chemically precipitated or recrystallized [61]. The grains may be oriented by the process of deposition or by later deformation [61]. For example, mica flakes are oriented as they are deposited. Nonclastic rocks can be crystalline, meaning that they are made up of interlocking crystals in contact with each other over the whole surface, or granuloid where the crystals show only tangential contact [44]. These crystals can have a variety of shapes and sizes and may have smooth or jagged boundaries.

Polish or gloss is related to the regularity of reflection and is indicated by the presence of highlights. It is produced mechanically by wear, particularly if the abrasive agent is of fine grain, or by the deposition of glaze [44]. However, most pebbles are dull.

Surface markings may include striations, scratches, percussion marks, indentations and pits [44]. Striations are produced from ice action, generally glacial ice. They are narrow, straight scratches clearly cut below the surface of the rock. Percussion marks, or crescentic impact scars, are caused by blows on the pebble surface. They are indicative of high-velocity flow. Surface indentations and pits are produced by etching and differential solution of inhomogeneities of a rock or may also be due to abrasion from adjacent pebbles.

Abrasion and transport tend to modify surface texture more easily than they modify other properties such as shape [44]. As such, surface texture, and surface markings in particular, typically reflect the last environment of deposition.

Sample rocks displaying some of these properties are shown in Figure 4. These characteristics are important to rock texture from a geology sense. However, there are more general approaches to visual texture analysis.

First of all, what is texture in an image? Looking up close at a leaf, it might be the pattern of veins visible. Examining a tree, it might be the structure of branches. When viewing satellite imagery, it is the lakes, rivers, forests, fields and hills. Evidently, texture depends on scale. It is fairly easy for us to look at an image and have an idea of what we recognize as texture. However, it is much harder to formally define it.

There are different types of textures. Structured textures consist of the same components repeated in a consistent pattern. This could be characters in a paragraph of text or a stack of identical soup cans. Stochastic textures, on the other hand, have a random element. There is a statistical pattern of elements, but the components are not exactly repeated. Examples of this case are a grassy field, a piece of wood or a rock. Structured textures are more easily handled because the identical components are recognizable and it is just their pattern of repeatability that must be represented. Stochastic textures are more difficult. Here we are dealing with rocks as a stochastic texture.

Recognizing similar textures and distinguishing different ones is not an easy task, even for humans. Illumination conditions and viewing angles can change the appearance of materials. This can

result in a single material appearing vastly different under different conditions, or various different materials appearing quite similar under certain conditions.

This section explores methods for characterizing rock texture. The fractal dimension, statistics of a gray-level co-occurrence matrix, directionality histograms and textons formed from responses of convolving the image with a filter bank are examined.

### 2.2.1 Related Work

An early approach to texture perception was that of Julesz [28, 27]. He developed numerous textures with different properties in an attempt to determine what characteristics make two textures distinguishable. His hypothesis that textures cannot be distinguished if they have identical second order statistics was later disproven, but his work provided many initial ideas for texture analysis.

One of the first representations developed for textures is that of fractals [29]. This applies best to structured textures, but can also be used on stochastic ones. Fractals display self-similarity on all scales. They need not show exactly the same structure on all scales, but the same type of structures must be present on all scales. The fractal dimension measures the degree of self-similarity. Natural fractals do not have same degree of self-similarity, but they do exhibit statistical self-similarity. This can be measured with the box dimension [29]. The fractal or box dimension gives a measure of roughness of a surface. Intuitively, the larger the fractal dimension, the rougher the texture is.

A second method for texture representation involves co-occurrence, the second-order statistics of images - that is, the spatial relationships between pairs of gray values of pixels [24, 53]. A co-occurrence matrix represents how often pairs of gray levels of pixels, that are a certain distance apart at a given angle, occur in an image. These statistics can be summed up in measures such as energy, entropy, contrast, homogeneity and correlation. As with fractals, this involves a significant loss of information, but hopefully reduces to just the textural properties that are important.

Texture representation with the use of textons is the most common approach taken over the last decade or so. Most often this is done by convolving the texture with a set of filters. A Gabor filter bank can pick out the frequency of information in an image. A bank using edges, bars and spots can identify other types of features present in a texture, similar to how it is thought the human brain processes some visual information [37, 54]. Different scales and orientations of these filters are usually used. By clustering the filter responses, a set of textons can be found as a basis for the representation of small structures in an image. From this, each pixel in an image can be assigned to the closest texton. The distribution of textons over an image region can now be used for segmentation, or in the form of a histogram over the whole image to provide a model for a texture.

A similar approach, also using textons, involves a Markov Random Field (MRF) [55]. The MRF assumption is that the probability of the intensity of a central pixel is dependent only on its local neighborhood and not on any other pixels in the image. The intensities of neighborhood pixels for all possible central pixels can then be aggregated and clustered to form textons. A model for a texture then takes the form of a 2D histogram of the textons over the image conditioned on the central pixel intensity. An important feature of this algorithm is that the neighborhood size used is much smaller than the size of the largest filter in the filter bank approach. This is useful when there is limited data available or when computing texture near boundaries in an image for segmentation.

Each of these methods for representing texture can be used for classification or segmentation, in addition to other applications like texture synthesis and shape from texture.

Now we turn to the subject of texture analysis applied to the specific case of rock geology. The most common method is that of the co-occurrence matrix [59, 43, 34]. Statistics like energy, entropy, contrast and many others can summarize important characteristics of rocks. An approximation of the fractal dimension [7] has also been used as a measure of rock texture [50]. As directionality is an important property of rocks, the directional histogram involving convolving the image with a set of directional masks at different orientations has provided useful texture information [34]. Finally, the

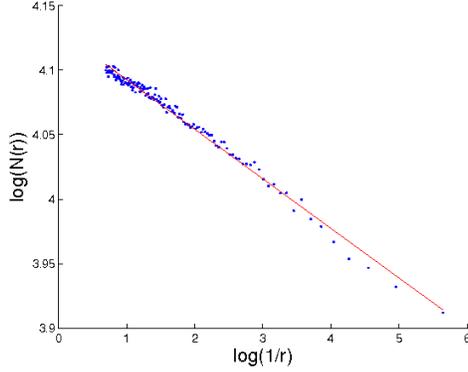


Figure 5: Points and best fit line for computing the fractal dimension for a texture.

use of Gabor filter banks can extract local scale, orientation and frequency information from textures [35, 36, 5, 23]. One thing that makes rock textures unique from other common homogeneous texture analysis is that they often consist of multiple textures. One approach to deal with this is to divide the image up into sub-images and characterize each separately [43]. Many of these methods will be explored further in this section.

## 2.2.2 Methods

This section details the methods chosen to represent texture. The specifics of each method are given and later experimental analysis will discuss the importance of each.

### 2.2.2.1 Fractal Dimension

The first approach taken for texture analysis is the use of the fractal dimension to characterize roughness as is used by Thompson [50]. The fractal dimension  $D$  of a set  $S$  is defined according to

$$1 = N_r r^D$$

where  $S$  is the union of  $N_r$  distinct copies of itself, each scaled down by a ratio  $r$ . The fractal dimension characterizes the degree of self-similarity of the texture. Intuitively, the greater the fractal dimension, the rougher the texture. To estimate the fractal dimension, we use in the Differential Box Counting method [7]. Take an image of size  $M \times M$ . If it is scaled down to a size  $s \times s$  then the ratio is  $r = s/M$ . For each value of  $s$ , draw a box of size  $s \times s$  on the image and determine the minimum and maximum gray level intensities,  $k$  and  $l$ , respectively, in the box. Then  $n_r(i) = l - k + 1$  is the  $i^{\text{th}}$  contribution for  $r$ . Doing this for multiple box sizes, we can compute  $N_r$  as

$$N_r = \sum_i n_r(i)$$

The boxes are chosen randomly throughout the image. When this is done for multiple values of  $r$ , the fractal dimension  $D$  is estimated from the slope of the best fit line in a  $\log(N_r)$  vs.  $\log(1/r)$  plot. This is shown in Figure 5.

As different textures may have the same fractal dimension, we compute the fractal dimension of five different images: the original image, a high gray-valued image, a low gray-valued image, a horizontally smoothed image and a vertically smoothed image [7]. The low and high gray-valued images are formed according to

$$I_{high}(i, j) = \begin{cases} I(i, j) - L_1 & \text{if } I(i, j) > L_1 \\ 0 & \text{otherwise} \end{cases}$$

$$I_{low}(i, j) = \begin{cases} 1 - L_2 & \text{if } I(i, j) > 1 - L_2 \\ I(i, j) & \text{otherwise} \end{cases}$$

where  $I$  is the original image,  $L_1 = g_{min} + g_{avg}/2$ ,  $L_2 = g_{max} - g_{avg}/2$ , and  $g_{max}$ ,  $g_{min}$  and  $g_{avg}$  are the maximum, minimum and average gray values in  $I$ , respectively. The smoothed images are formed according to

$$I_{horizontal}(i, j) = \frac{1}{2w + 1} \sum_{k=-w}^w I(i, j + k)$$

$$I_{vertical}(i, j) = \frac{1}{2w + 1} \sum_{k=-w}^w I(i + k, j)$$

for  $w = 3$ . The fractal dimension is computed for each of these five images to form a vector of features for the texture.

### 2.2.2.2 Co-occurrence Statistics

The Gray-Level Co-occurrence Matrix (GLCM) measures spatial relationships of pixels in an image. The matrix is defined as

$$GLCM_{d,\alpha}(i, j) = |\{(r, s), (t, v) : I(r, s) = i, I(t, v) = j\}|$$

where  $d$  is the distance at an angle  $\alpha$  between pixels of intensities  $i$  and  $j$  and  $|\cdot|$  is the cardinality of a set. In other words, entry  $(i, j)$  is the number of occurrences of the pair of gray levels  $i$  and  $j$  at a distance  $d$  and angle  $\alpha$  apart. This is computed for  $\alpha = 0^\circ, 45^\circ, 90^\circ, 135^\circ$  and  $d = 1$  to 5, averaged over these values for pixel intensities divided into eight bins. From this, the following features can be computed:

$$Contrast = \sum_{(i,j)} |i - j|^2 GLCM(i, j)$$

$$Correlation = \sum_{(i,j)} \frac{(i - \mu_i)(j - \mu_j) GLCM(i, j)}{\sigma_i \sigma_j}$$

$$Energy = \sum_{(i,j)} GLCM(i, j)^2$$

$$Homogeneity = \sum_{(i,j)} \frac{GLCM(i, j)}{1 + |i - j|}$$

These features are computed by standard Matlab functions and form a vector to represent the texture by. This method is similar to that used by Partio [43] and Lepisto [34].

### 2.2.2.3 Directional Histogram

As directionality is an important geologic feature of rocks, we develop a method to characterize this with a directional histogram [34]. A set of directional masks, shown in Figure 6, is convolved with the image, with the result summed over all pixels giving the contribution of that mask direction in the image. This is done for each mask and the histogram is normalized to sum to one. As it is not important which direction is dominant in the texture, only the distribution of directions, we rotate the elements of the histogram so that the largest value is in the first element.

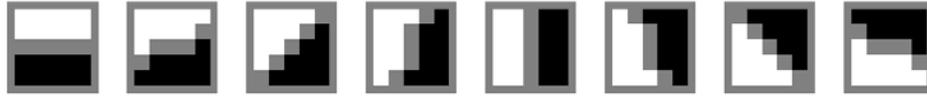


Figure 6: Directional masks used to compute directional histograms.

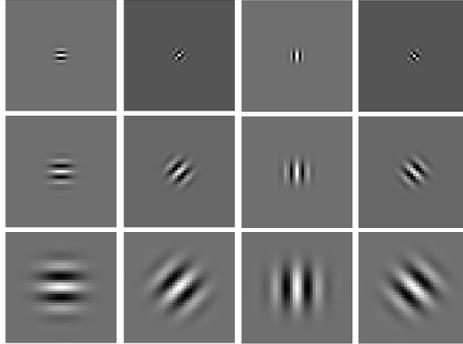


Figure 7: Gabor filter bank.

#### 2.2.2.4 Textons

The most common approach to texture analysis in computer vision involves convolving the texture with a set of filters and clustering the responses to form textons. Here we try two different filter banks: the Gabor filter bank and the Maximum Response 8 (MR8) filter bank.

A Gabor function is a 2-D Gaussian modulated with a sinusoid. A Gabor filter bank consists of a series of Gabor filters at various scales and orientations. The filter bank used here is shown in Figure 7. We use four orientations and three scales, resulting in a 12-dimensional response vector for each pixel [5].

We also use Varma and Zisserman's MR8 filter bank [54] shown in Figure 8. Here there is an edge filter (first derivative of a Gaussian) and a bar filter (second derivative of a Gaussian) at six orientations and three scales, and two spot filters (a Gaussian and a Laplacian of a Gaussian). The unique trick with this approach is that for each filter, the maximum response is taken across the six orientations. This reduces the response vector down to eight dimensions. By doing this, potentially useful data is lost; however, a smaller dimensional vector can be useful for reducing the quantity of data in later processing.

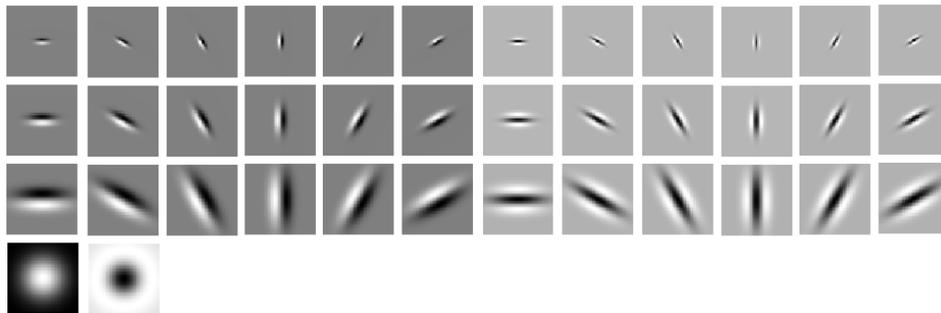


Figure 8: MR8 filter bank.

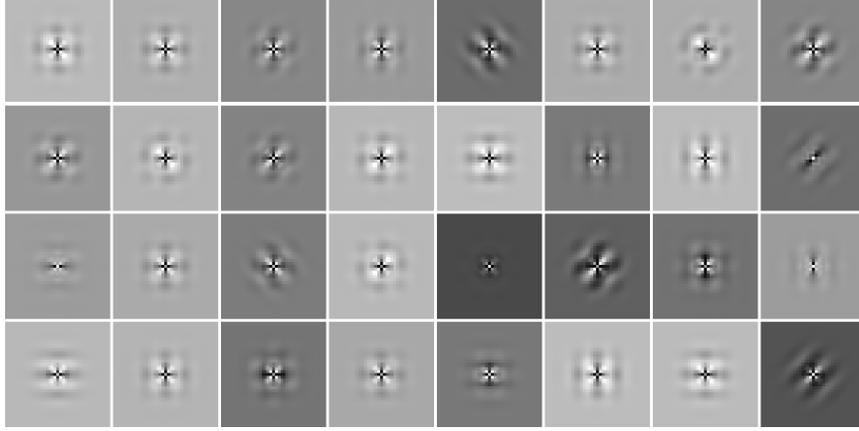


Figure 9: Textons formed with the Gabor filter bank.

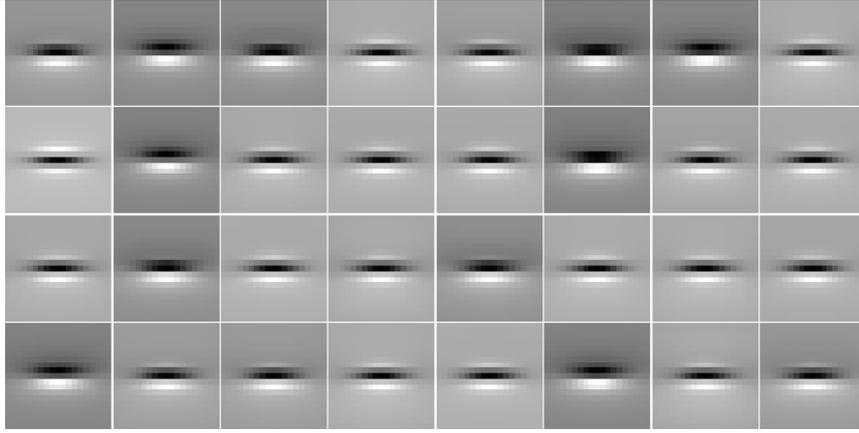


Figure 10: Textons formed with the MR8 filter bank.

Each response vector is normalized according to

$$R_i \leftarrow R_i \frac{\log(1 + \|R\|/0.03)}{\|R\|}$$

where  $R$  is the vector and  $R_i$  is each element.

Once the filter bank is convolved with the texture to form a response vector for each pixel using either the Gabor or MR8 filter bank, we then compute a set of textons. The response vectors from all textures in the set are aggregated into one long matrix of size  $M \times N$  where  $M$  is the total number of pixels in the images and  $N$  is the dimensionality of the response vector. The response vectors are clustered using k-means. Here we use 32 clusters, with each cluster representing a texton. A sample set of textons is shown in Figure 9 for the Gabor filter bank and in Figure 10 for the MR8 filter bank. Each texton is displayed as the linear combination of the filters and represents a form of primitive structure in the set of textured images. For the MR8 filter bank, all textons are horizontal because we are only showing one orientation.

From this we can compute the closest texton for the response vector at each image pixel, forming a texton map. When dealing with rocks, only the pixels contained within the rock are used for

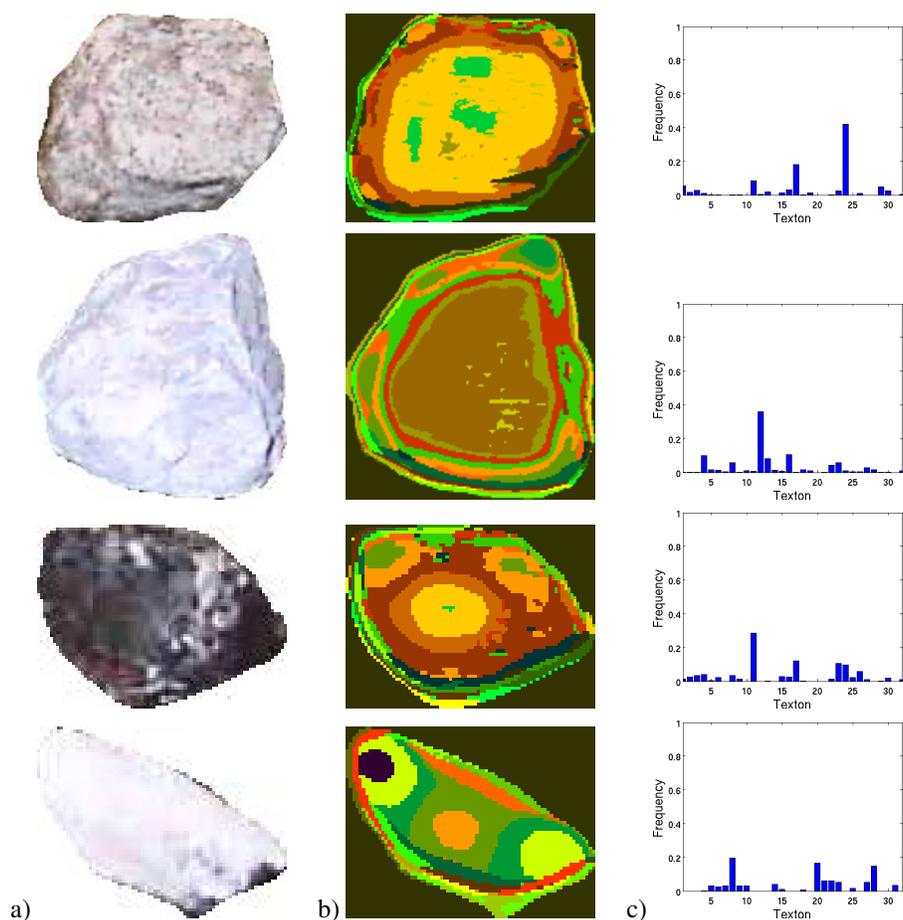


Figure 11: a) Sample rocks, b) their texton maps using MR8 filters showing the maximally-responding texton at each pixel (each color represents a different texton), and c) their texton histograms.

forming textons and the texton map. Sample rocks, their texton maps and texton histograms are shown in Figure 11. This distribution of textons can be represented with a histogram, counting the number of occurrences of each texton in the image. Textures can now be compared using the  $\chi^2$  statistic of their representative histograms.

### 2.2.2.5 Summary

The four texture analysis methods described here characterize texture in different ways. An overview is provided in Table 2 to summarize each method. It should be noted that in the computer vision community the fractal dimension and co-occurrence methods are typically viewed as out-dated. However, these are the methods that have commonly been applied to rock texture analysis, with texton approaches only appearing in the literature very recently. All methods were described for completeness and in Section 5 their effectiveness will be compared.

Texture analysis will later be used in rock detection (Section 3) and in geologic classification (Section 4). For the purpose of segmenting rocks from the background, Figure 31 provides an indication of the degree to which rock textures stand out from that of the surrounding material. It will

Table 2: Summary of texture analysis methods.

Method	Description
Fractal Dimension	<ul style="list-style-type: none"> <li>• Determine minimum and maximum gray level in numerous sub-image boxes of various sizes.</li> <li>• Sum contribution from each box at each scale <math>r</math>, giving <math>N_r</math>.</li> <li>• Fractal dimension is the slope in a <math>\log(N_r)</math> by <math>\log(1/r)</math> plot.</li> </ul>
Co-occurrence Statistics	<ul style="list-style-type: none"> <li>• Gray-Level Co-occurrence Matrix measures spatial relationships of pixels.</li> <li>• Compute contrast, correlation, energy and homogeneity from this matrix.</li> </ul>
Directional Histogram	<ul style="list-style-type: none"> <li>• Convolve with directional masks.</li> <li>• For each mask, sum responses over image.</li> <li>• Form normalized histogram.</li> </ul>
Textons	<ul style="list-style-type: none"> <li>• Use a Gabor or MR8 filter bank.</li> <li>• Convolve filters with all images.</li> <li>• Cluster responses to form textons.</li> <li>• Compute nearest texton for each image pixel.</li> <li>• Form texton histogram.</li> </ul>

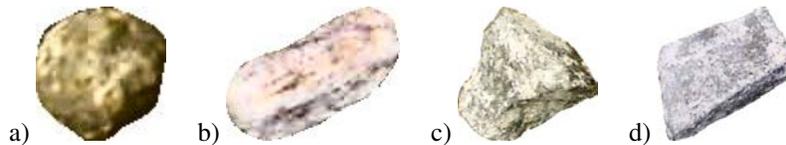


Figure 12: Sample rocks showing the variations in shape: a) spherical, b) elliptical and very round, c) angular d) straight-edged, but sharp corners.

later be shown in Section 5.3 that this measure of texture is usually sufficient for distinguishing rocks from the background.

## 2.3 Shape

When examining rocks and sediments found in the field, geologists find shape to be an important property to characterize. Shape properties can be divided into two categories: form and roundness. Form characterizes the overall shape of a rock, such as how close to spherical it is. Roundness measures the extent to which the edges have been smoothed out. Figure 12 illustrates variations in form and roundness that a rock may exhibit.

Rock form mainly reflects the conditions of deposition. It depends on the properties of minerals in the parent rock and their fracture habits [61]. It has a limited affect by abrasion [44]. Form has a major influence on hydrodynamical behavior as it affects the settling velocity and mode of transport.

A quantitative form measure requires a standard of reference. The sphere is most often chosen

for this purpose, as in the case of the metric sphericity. A sphere is approximately the limiting shape upon prolonged abrasion. It also has the least surface area for a given volume, resulting in the greatest settling velocity in a fluid of any possible shape [44]. Thus, more spherical particles tend to be deposited while others of the same size and density are carried away.

Roundness is caused by impacts with other grains during transport. It increases with distance of travel, most rapidly first, then slowly levels off [30, 44]. Experimental work and field studies show that using roundness, the distance of transport can be estimated to the right order of magnitude [44].

Abrasion also increases sphericity, but more slowly than roundness. Some particles, such as disc-shaped ones, may never become spherical. For example, mica is never found in a sphere [57]. The ultimate shape a pebble reaches is dependent on abrasion during transport and its initial shape [44].

Although rock sphericity, roundness and size are distinct properties, geology experiments have shown that there is a correlation between them. Larger rocks of a given material tend to be better rounded than smaller ones. Larger grains are transported mostly by rolling and sliding on the stream bottom, resulting in a rapid loss of sharp corners. Smaller grains can travel along the stream bottom or in the moving fluid with less contact with other particles resulting in a slower rounding process [2]. Larger grains also impact with more force. The same correlation is true for sphericity, but only in materials that have undergone appreciable wear [44]. Thus, a mature sediment will show a clear relation between size, sphericity and roundness, while in an immature one the relationship is much less pronounced [44].

Some rock shapes are distinctive of certain processes. One example is round stones that have undergone breakage, leaving most of the particle very round, with a sharp corner and roughness down one side. This breakage is indicative of exceptionally-high velocity currents [44].

Sphericity and roundness are most commonly used by geologists as metrics to describe the shape of a rock. This section explores these and other possible shape descriptors that characterize properties of rocks and might be useful features in geologic classification. A method for computing each measure is outlined.

### 2.3.1 Related Work

Sphericity and roundness are the most common rock shape descriptors employed. Sphericity is how close to spherical a rock is, while roundness measures the smoothness of its edges. These have both been used in geology since at least the 1930s, with their definitions evolving over time [1].

Both sphericity and roundness are very tedious to calculate. Thus, a chart was developed by Powers to allow the estimation of these metrics by visual inspection [46]. Similar versions of this chart have been created with different numbers of categories and different scales for sphericity and roundness [13, 44, 31]. Typically, the chart has a roundness scale down the side and sphericity across the top. Each cell displays a rock in that category. A chart by Crofts [13] is shown in Figure 13.

Graphs involving the ratio of the principal axes of the rock have also been used for categorization. The Zingg shape classes divide rocks into oblate, equant, bladed and prolate. These are useful because some definitions of sphericity fail to distinguish between disc-shaped rocks which are elongated in two directions and roller-shaped rocks which are elongated in one direction [44].

Numerous other shape measures have been proposed, but none are as commonly used as sphericity and roundness. Some have been used to characterize rocks [1, 15, 17, 16, 22], while others have been used for powder or sand grains but not larger rock sediments [3, 25, 32, 39]. Many have been used in shape analysis in computer vision, but have not commonly been applied to rock geology [14, 45]. Some measures describe the geometrical shape using properties such as convexity or by comparing the outline to a circle or ellipse [14, 39, 45]. Others find ways to approximate the fractal dimension of a shape [14, 25]. The most common recent technique is to transform the shape edge points into a one-dimensional signal by computing their distance from the centroid and applying

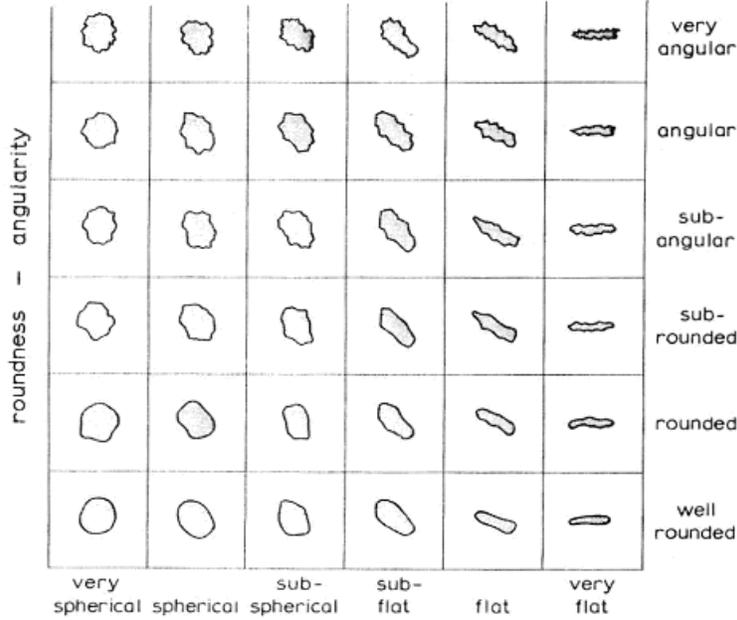


Figure 13: Crofts' sphericity and roundness chart [13].

Fourier analysis [3, 14, 15, 17, 19]. For a more complete analysis of rock shape, Oakey [42] develops a method using the distribution of form properties in a large sample of rocks, while Durian [18] examines the statistical distribution of the curvature along rock edges. Many of these techniques are explored in this section.

### 2.3.2 Methods

To accurately characterize rock shape, a three-dimensional representation is necessary. However, to simplify the task, here we work in two-dimensions with an image of a rock. If the selected viewpoint of the rock is representative of the shape of the rock, this approach should be sufficient. However, if a very elliptical rock is viewed from one end or a rock is round on one side but not the other due to a fracture, this assumption will break down.

We work from a binary mask showing the silhouette of the rock. Our purpose is to compute rock shape metrics. These are divided into the categories of those describing form, those describing roundness and those that do not clearly belong to the form or roundness categories. This section examines the details of the methods explored.

Some approaches to object shape recognition involve comparing the outline of the object, or the output from an edge detector, to a known template. This could be used for rock classification by using a training set of multiple rocks in each class, and applying a nearest neighbor classifier to each novel rock image with a distance measure such as the Chamfer distance [13]. However, this method does not scale or generalize well as templates are required for each class. From a geology perspective, it is likely that rocks are better characterized by specific shape properties, rather than templates for each class, as there are many different possible shapes that a rock can take on. This is especially the case for roundness shape properties. This report does not describe these types of methods any further.

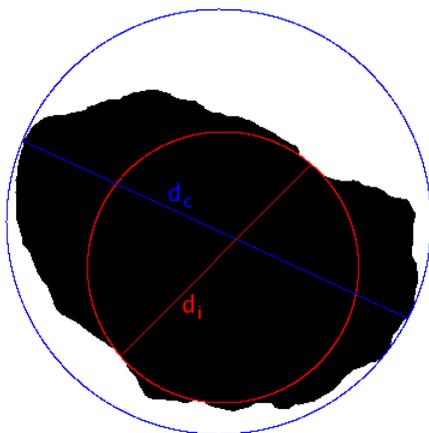


Figure 14: Minimum circumscribing and maximum inscribed circles of a rock for computing Riley sphericity.

### 2.3.2.1 Form

Form describes the overall shape of a rock, such as how close to a sphere it is. The following measures can be used in this task:

**Riley Sphericity** This measure was originally defined by Wadell as

$$WadellSphericity_1 = \frac{s}{S}$$

where  $s$  is the surface area of a sphere of the same volume as the rock and  $S$  is the actual surface area of the rock [56]. He later modified this to

$$WadellSphericity_2 = \sqrt[3]{\frac{d_I d_S}{d_L^2}}$$

where  $d_S$ ,  $d_I$  and  $d_L$  are the principle axes of the rock [8]. Both of these definitions require knowledge of the 3-dimensional structure of the rock. Riley later suggested a calculation that only involves measurements from a 2-dimensional view. He defined sphericity as

$$RileySphericity = \sqrt{\frac{d_i}{d_c}}$$

where  $d_i$  is the diameter of the maximum inscribed circle and  $d_c$  is the diameter of the minimum circumscribing circle [47]. The inscribing and circumscribing circles are shown in Figure 14. The value of sphericity ranges from 0 to 1, with a perfect circle having a value of 1.

To compute the Riley sphericity, it is first necessary to find the maximum inscribed and minimum circumscribing circles of the rock. One method for this involves the convolution of circles of various sizes with the mask image. The convolutions are done by multiplying in the Fourier domain for computational efficiency. When the mask image is convolved with a circle of a given radius, the maximum value of the convolution will be equal to the area of the rock only if the circle circumscribes the rock. Also, the maximum value of the convolution will be equal to the area of the circle only if the circle fits inside the rock. Thus, to determine the radius of the inscribed circle, we

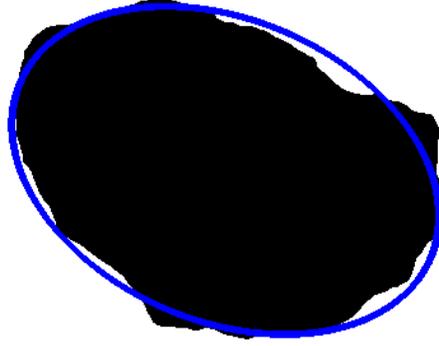


Figure 15: Best fit ellipse for computing elongation.

start with a circle of radius one and convolve with the image, progressively making the circle larger. Once the maximum value of the convolution becomes less than the area of the circle, the circle no longer fits inside the rock and we have found the desired radius. Next, to determine the radius of the circumscribing circle, we start with a circle of radius equal to the inscribed circle radius and progressively make the circle larger. Once the maximum value of the convolution becomes equal to the area of the rock, the circle circumscribes the rock and we have the required radius.

More efficient algorithms for computing the necessary circle diameters are available [49]; however, the method used here is simple to implement and computation time is not currently an issue.

**Elongation** The ratio of the minor and major axes of the best-fitting ellipse forms the metric elongation [45]. The ellipse is computed by minimizing the algebraic distance of each edge point to a quadratic curve with the added constraint that the quadratic is an ellipse [52]. An example of such an ellipse is shown in Figure 15.

**Ellipse Error** By comparing the region to its best fit ellipse we can compute the ellipse error. This is done by averaging the distance from each rock boundary point to the closest point on the ellipse. The distance between the data point and the ellipse is approximated as the distance between the data point and the point on the ellipse that lies on the line connecting the data point and the ellipse center [22].

**Circular/elliptic variance** Here we assume that a standard rock closely resembles a circle or an ellipse and measure the deviation of the edge points from that standard shape. These measures are defined as

$$CircularVariance = \frac{1}{N} \sum_{i=1}^N (\|p_i - \mu\| - \mu_r)^2$$

and

$$EllipticVariance = \frac{1}{N\mu_{rc}} \sum_{i=1}^N \left( \sqrt{(p_i - \mu)^T C^{-1} (p_i - \mu)} - \mu_{rc} \right)^2$$

where

$$\mu = \sum_{i=1}^N p_i, \mu_r = \sum_{i=1}^N \|p_i - \mu\|, \mu_{rc} = \sum_{i=1}^N \sqrt{(p_i - \mu)^T C^{-1} (p_i - \mu)}, C = \frac{1}{N} \sum_{i=1}^N (p_i - \mu)(p_i - \mu)^T,$$

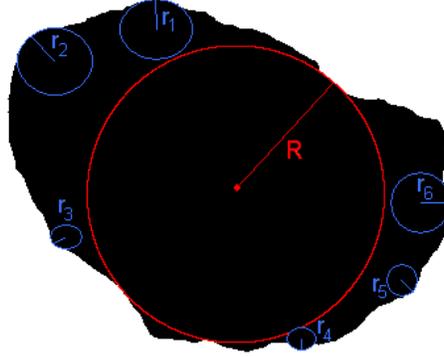


Figure 16: Radius of curvature of some corners of a rock and maximum inscribed circle.

$p_i = (x_i, y_i)$  is the  $i^{th}$  contour point and  $N$  is the number of contour points [45].

### 2.3.2.2 Roundness

Roundness measures the extent to which the edges and corners of a rock have been smoothed or rounded. Metrics for this include the following:

**Wadell Roundness** Wadell defined roundness as

$$WadellRoundness = \frac{\sum_{i=1}^N r_i}{NR}$$

where  $r_i$  are the radii of curvature of the corners of the rock,  $N$  is the number of corners and  $R$  is the radius of the maximum inscribed circle [56]. This is shown for some of the corners of a rock in Figure 16. By Wadell's definition, a corner is any point on the rock contour such that the radius of curvature is equal to or less than the radius of the maximum inscribed circle. Ideally, roundness should be measured in three orthogonal planes, but for most rocks a single plane is sufficient [56]. Roundness approaches a value of one for perfectly round rocks.

In order to compute Wadell roundness, we need to first determine the radius of curvature for all edge points on the rock. One method of doing this is to determine the radius of the largest circle that is tangent to the rock at that point, and is contained within the rock. This is done by first creating a square matrix with ones representing the location of a circle of radius  $R = 1$ , and zeros outside. For each possible pixel location of the center of this circle such that it is tangent to the rock edge at the given point, perform element-by-element multiplication of the circle with this part of the mask image and sum these numbers. If this sum is equal to the area of the circle, then the circle fits inside the rock and we repeat this procedure, incrementing the radius by one. If no locations are found for which the circle fits inside the rock, then the process is terminated and the radius of curvature is  $R - 1$ .

This computation is quite expensive, especially for larger rocks with larger radii of curvature. Optimization could be done by using the inward normal at the edge point, determined from the image gradient, as an indication of where the center of a tangent circle is likely to be. However, we would still need to search within  $\pm 45^\circ$  of this direction to obtain an accurate result.

Once the radius of curvature is known for each edge point, we simply take the average of those that are smaller than the radius of the maximum inscribed circle. Dividing this average by the radius of the maximum inscribed circle gives the required result.

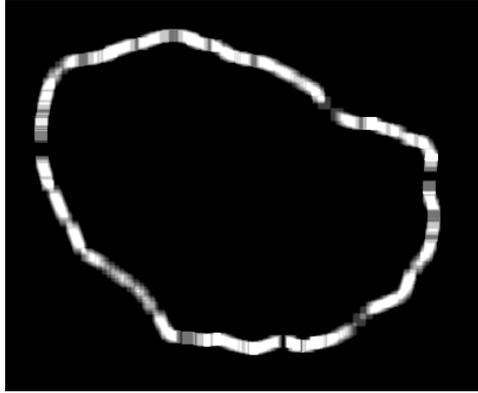


Figure 17: Corner strength where white indicates a stronger corner.

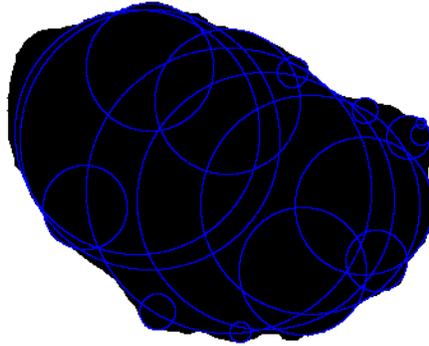


Figure 18: Radius of curvature of the corners of a rock.

Although Wadell states that a corner is any point on the rock contour such that the radius of curvature is smaller than the radius of the maximum inscribed circle, other suggested definitions use only the sharpest corner or sharpest two corners [26, 1]. This suggests that other possibilities are to use only the largest corners or only the strongest corners with a measure of corner strength.

A method for locating the corners of a rock is the algorithm developed by Tomasi and Kanade [52]. This is performed for each edge point, obtaining a numerical value indicating its strength as a corner. The corner strengths of a rock are shown in Figure 17. By inspection, a neighborhood size of 11x11 pixels produces good results. Next, the corner with the greatest strength is located and the strength for all edge points within the neighborhood used by the corner detector are discarded. This process of selecting the strongest edge is repeated until the strength of all edge points have been discarded and we have our selection of the corners for the rock. The radius of curvature of the corners of a rock are shown in Figure 18.

**Angularity** This method computes a measure of curvature for boundary points of the region and takes the standard deviation [22]. The boundary points are labeled in order around the region and we wish to estimate the curvature of a point  $P_i$ . This is done by fitting triangles with the apex resting at  $P_i$  and the remaining two vertices,  $P_h$  and  $P_j$ , on either side of  $P_i$  at various distances. We

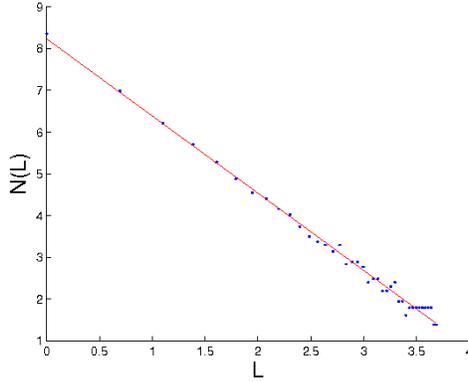


Figure 19: Points and best-fit line for computing the box dimension.

place  $P_h$  and  $P_j$  subject to

$$\begin{aligned} h < i - \alpha & \quad j > i + \alpha \\ h > i - \beta & \quad j < i + \beta \end{aligned}$$

The variables  $\alpha$  and  $\beta$  are set by experimentally determining the values which produce the most useful results. All possible combinations of  $h$  and  $j$ , subject to the above constraints, are used to compute the angle  $\angle hij$  as

$$\angle hij = \arccos \frac{a^2 + b^2 - c^2}{2ab}$$

with  $a = \|P_h - P_i\|$ ,  $b = \|P_j - P_i\|$  and  $c = \|P_h - P_j\|$ . The estimated curvature is then characterized as the maximum  $\angle hij$  for all possible  $h$  and  $j$ . Finally, the angularity is represented as the standard deviation of the estimated curvature at all boundary points.

**Fractal Dimension** The roundness of a rock can be represented by the fractal dimension. The first method of estimating this is with the box dimension [14]. The rock mask image is divided up into boxes of size  $L$  and the number of non-empty boxes is computed as  $N(L)$ . The box dimension is defined according to the equation

$$N(L) = \left[ \frac{L_{max}}{L} \right]^D$$

where  $L_{max}$  is the maximum dimension of the rock mask and  $D$  is the box dimension. Rearranging, this equation takes the form

$$\log(N(L)) = -D \log(L) + D \log(L_{max})$$

Thus, the box dimension,  $D$ , can be estimated by computing  $N(L)$  for numerous values of  $L$  and estimating the slope of the best fit line in a  $\log(N(L))$  by  $\log(L)$  plot. This is shown in Figure 19.

Another way of measuring the fractal dimension is with the divider method [25]. The rock mask image is divided up by horizontal lines at a distance  $\lambda$  apart, as shown in Figure 20. The intersection of the lines with the rock outline are found. By connecting these points, a new shape similar to that of the rock is formed and the perimeter is computed. When this is done for multiple values of  $\lambda$ , the perimeter and  $\lambda$  values have an approximately linear relationship shown in Figure 21. The slope of the best-fit line approximates the fractal dimension.

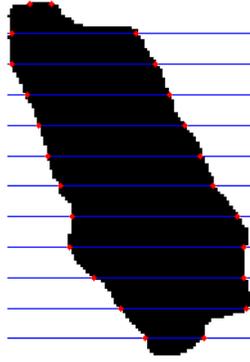


Figure 20: Divider method for approximating the fractal dimension.

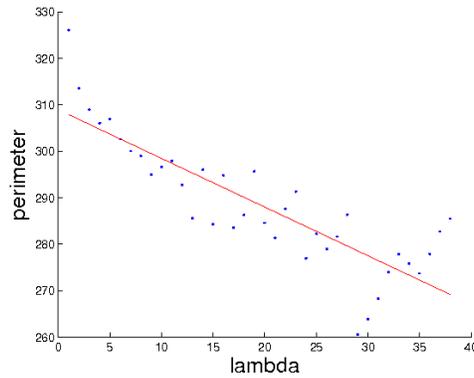


Figure 21: Points and best-fit line for computing the fractal dimension by the divider method.

**Diepenbroek Roundness** Diepenbroek developed a method of characterizing the roundness of a rock using Fourier analysis [15]. This method is intended to produce similar results to Wadell's definition of roundness. First, the edge points of the rock are transformed into polar co-ordinates  $(r, \theta)$  with  $r$  representing the distance from the rock centroid and  $\theta$  as the angle. Using linear interpolation, the distance from the centroid is determined for 256 equally spaced angles. Fast Fourier Transform is performed on this signal. The amplitude for the lower harmonics characterizes the overall shape of the rock, while the higher harmonics often characterize noise due to digitization. Thus, lower harmonics are discarded as they characterize form rather than roundness, and higher harmonics are omitted due to noise. The remaining harmonics are used in the rest of this analysis. The rock edge points in polar co-ordinates and the Fourier transform are shown for a rock in Figure 22.

Simply summing the amplitude of all harmonics would not give enough weight to the smaller details of the rock contour. Thus, a normalization function must be found. This is done by computing the average amplitude at each harmonic for the rocks in the data set. An exponential function is found to fit the average amplitude. The normalization function chosen is

$$c_k = 0.49 \exp(0.17k)$$

where  $k$  is the harmonic number. This is shown in Figure 23.

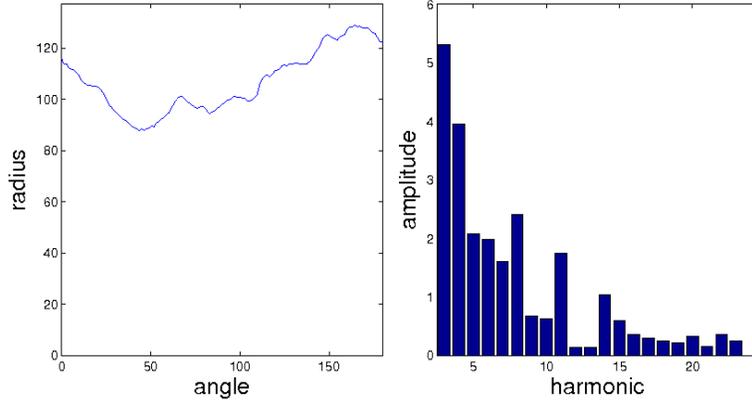


Figure 22: Rock contour in polar co-ordinates and its Fourier transform.

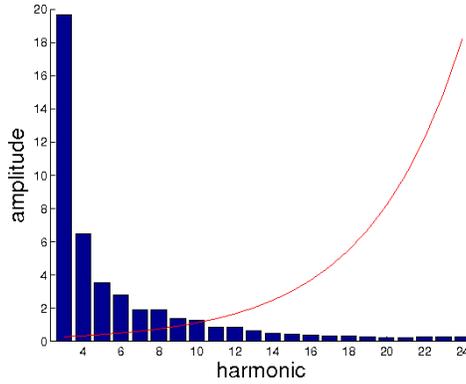


Figure 23: Mean Fourier transform (blue) and normalization function (red).

A preliminary value for roundness can now be computed as

$$DiepenbroekRoundness_{pre} = \sum_{i=n_1}^{n_2} c_k A_k$$

where  $A_k$  is the amplitude of the  $k^{th}$  harmonic, and  $n_1$  and  $n_2$  are chosen to achieve the best performance with the metric.

In Diepenbroek's work, the effect on the spectrum of the best-fitting ellipse is also removed [15]. However, in experiments it was found that the metric performs best without this extra step. Some of the steps described here are slightly modified from those by Diepenbroek [15] as the constants in the equations must be determined to fit the particular data set of rocks used.

### 2.3.2.3 Other Metrics

In this section we provide details for rock shape metrics that do not clearly fall into the categories of form or roundness. Form metrics involve comparing the overall shape of the rock to a standard shape, such as a circle or an ellipse. Roundness metrics involve some notion of the sharpness of the corners of the rock or, in the case of the fractal dimension, a measurement of the roughness of the rock boundary. The metrics in this section characterize other properties of rocks, such as convexity.

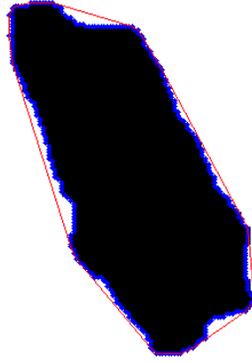


Figure 24: Convex hull of a rock for computing convexity.

**Compactness** This metric is defined as

$$Compactness = \frac{\sqrt{Area}}{Perimeter} [45].$$

The area and perimeter of a region can be computed with standard Matlab functions.

**Circularity** This metric is defined as

$$Circularity = \frac{Perimeter^2}{Area} [39].$$

**Convexity** By comparing the region of the rock to its convex hull we can compute convexity [45]. The convex hull of a rock is displayed with the red outline in Figure 24. Convexity can then be characterized in two ways:

$$ConvexityPerim = \frac{Perimeter_{convexhull}}{Perimeter_{rock}}$$

or

$$ConvexityArea = \frac{Area_{convexhull}}{Area_{rock}}.$$

Computing the convex hull, perimeter and area of a region can be done with standard Matlab functions.

#### 2.3.2.4 Summary

In this section we have explored numerous metrics for characterizing rock shape. Table 3 provides an overview of these methods. Each characterizes shape in a different but unique way and their effectiveness will be explored in Section 5.

## 3 Rock Detection and Segmentation

This section explores the process of detecting and segmenting rocks from images. This task is usually straight-forward for humans: just look at an image, identify the rocks and draw a boundary around each. It may be a slow and tedious process, but generally achieves very good results. Like so

Table 3: Summary of shape characterization methods.

<b>Form</b>	
<b>Method</b>	<b>Description</b>
Riley Sphericity	$\sqrt{d_i/d_c}$ where $d_i$ is the diameter of the maximum inscribed circle and $d_c$ is the diameter of the minimum circumscribing circle.
Elongation	Ratio of axes of best-fit ellipse.
Ellipse Error	Average distance from each rock boundary point to the closest point on the best-fit ellipse.
Circular Variance	$\frac{1}{N} \sum_{i=1}^N (\ p_i - \mu\  - \mu_r)^2$
Elliptic Variance	$\frac{1}{N\mu_{rc}} \sum_{i=1}^N \left( \sqrt{(p_i - \mu)^T C^{-1} (p_i - \mu)} - \mu_{rc} \right)^2$
<b>Roundness</b>	
<b>Method</b>	<b>Description</b>
Wadell Roundness	$\sum_{i=1}^N r_i / NR$ where $r_i$ is the radius of curvature of a corner $i$ , $R$ is the radius of the largest inscribed circle and $N$ is the number of corners.
Angularity	Standard deviation of the curvature at all boundary points.
Fractal Dimension	Computed using the box or divider method.
Diepenbroek Roundness	<ul style="list-style-type: none"> <li>• Compute the distance of each boundary point from the centroid to form a 1-D signal.</li> <li>• Take a weighted sum of the Fourier transform of the signal.</li> </ul>
<b>Other</b>	
<b>Method</b>	<b>Description</b>
Compactness	$\sqrt{Area/Perimeter}$
Circularity	$Perimeter^2/Area$
Convexity Perimeter	$Perimeter_{convexhull}/Perimeter_{rock}$
Convexity Area	$Area_{convexhull}/Area_{rock}$

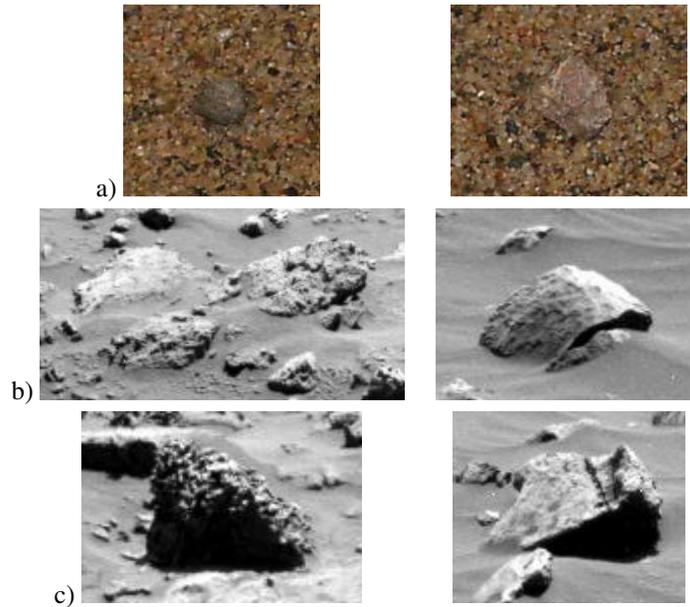


Figure 25: Examples of rocks that may be difficult to detect: a) rocks that blend in quite well with the background, b) rocks that are partially covered in sand, c) rocks with cast shadows due to directional lighting. (Images b) and c) courtesy of NASA/JPL-Caltech.)

many computer vision problems, this is a much more difficult task for a computer than for humans. However, in order to process large amounts of data, it must be automated.

There are numerous complications that can make rock detection and segmentation difficult. Illumination conditions such as low lighting or highly directional lighting causing cast shadows can complicate detection. Rocks can also overlap each other, be partially covered in sand or have a color and texture similar to the substrate. Figure 25 illustrates these cases.

Rocks come in all different colors, shapes and sizes. Here we will not attempt to define a rock, except to say that we are looking for any objects that a human would clearly identify as a rock protruding from the background material. To make this process feasible, we are attempting to find all rocks larger than 40 pixels in diameter. We wish to find as many rocks as possible in each image, while minimizing the number of false detections and accurately localizing the boundary of each rock. This boundary localization is important for further processing, such as analyzing the shape of each rock.

In this work we take the approach of first over-segmenting the image, and then using top-down knowledge about rocks to intelligently merge regions and identify rocks. The bottom-up approach to segmentation uses intensity, color, texture and boundary contours, each a low-level cue. Higher-level knowledge of rocks, such as their shape, shading due to illumination, boundary contours, color and texture, is necessary for accurate detection. By combining the bottom-up and top-down knowledge in this manner, we hope to use all available information to accurately segment images and detect rocks.

An important note in detecting the boundary of rocks is that we are not simply looking for edges. An edge is a change in intensity from one pixel to the next and can be easily identified with the Canny edge detector. What we are looking for are boundaries: contours in the image plane that represent a change from one object or surface to another. Boundaries are identified by differences in intensity or texture between regions. However, just boundaries are not sufficient as we require a closed contour

around the rock. Thus, more complicated segmentation algorithms are necessary.

### 3.1 Related Work

Some simple approaches to rock detection can perform the task well when certain simplifying assumptions are made. Crida [11] uses a multi-resolution system. Each pixel is hypothesized as the center of a fixed-size rock and features such as edge boundary strength, discontinuity of boundary and shape are computed. The features are thresholded to identify rocks. This is based on the assumption that rocks are basically circular. A similar approach searches for elliptical shaped rocks using the Hough transform [12]. This method assumes that rocks are lighter than the background, have a continuous boundary and a connected interior. Rocks are again classified using a similar feature set.

A more general approach removes many of the simplifying assumptions. In Thompson's work [50], each color channel in Hue Saturation Intensity is segmented separately. The image is shattered into 5x5 pixel regions which are iteratively rejoined when their mean pixel values are close enough. Rock detection is then performed using a Bayesian belief network operating on the following features: the ratio of squared perimeter to area, difference in the mean and variance of color between the region interior and a context region, texture measured by the fractal dimension, height above the ground plane, magnitude of the intensity gradient and absolute color. Computing these and other features also allows for the classification of detected rocks.

Yet another approach to rock detection involves segmenting rocks from the background using range data [22]. A plane is fitted to the ground and the height map is segmented to locate the rocks.

A common problem to all these methods is that they do not accurately localize the boundary of rocks. Depending on the application, this may not be an issue. However, for the task explored in this paper we are required to analyze the albedo, color, texture and shape of rocks. Particularly for shape measurements, accuracy in boundary localization is a necessity. Thus, we now examine more accurate methods for segmentation that may not have been applied to the explicit application of rock detection.

Numerous segmentation algorithms are available for different types of tasks. The simplest method is just to apply a threshold - all pixels greater than a given intensity belong to one segment, while those less than are another. This would work in a simple scene where all rocks are either darker or lighter than their background. However, most scenes are more complicated with variations in lighting within the background and individual rocks due to differences in reflectivity and illumination. More sophisticated algorithms that have performed well and are now commonly used are Blobworld [4], mean shift [10] and graph-based algorithms [20, 48]. Some of these will be examined in the application of rocks in the next section. A quality segmentation algorithm can help to obtain a pixel-accuracy boundary outline for a rock.

High-quality superpixel segmentations have previously been used to aid in object detection. Body limbs and torsos in a variety of configurations are found using this method [41]. After computing the superpixel segmentation, features are evaluated on each superpixel. The features contour, shape, shading and focus help to identify superpixels containing body parts and a search is done of nearby superpixels in an attempt to find full limbs. A similar approach is taken in this paper by using a superpixel segmentation to identify and merge superpixels to detect rocks.

### 3.2 Approach

The method chosen for rock detection and segmentation involves first computing a superpixel segmentation of the image at multiple scales. Two classifiers are trained using examples of rocks and non-rock superpixels. The first classifier using intensity, color and texture features is applied to find superpixels that are likely to be part of rocks. Next, all combinations of superpixels that passed the

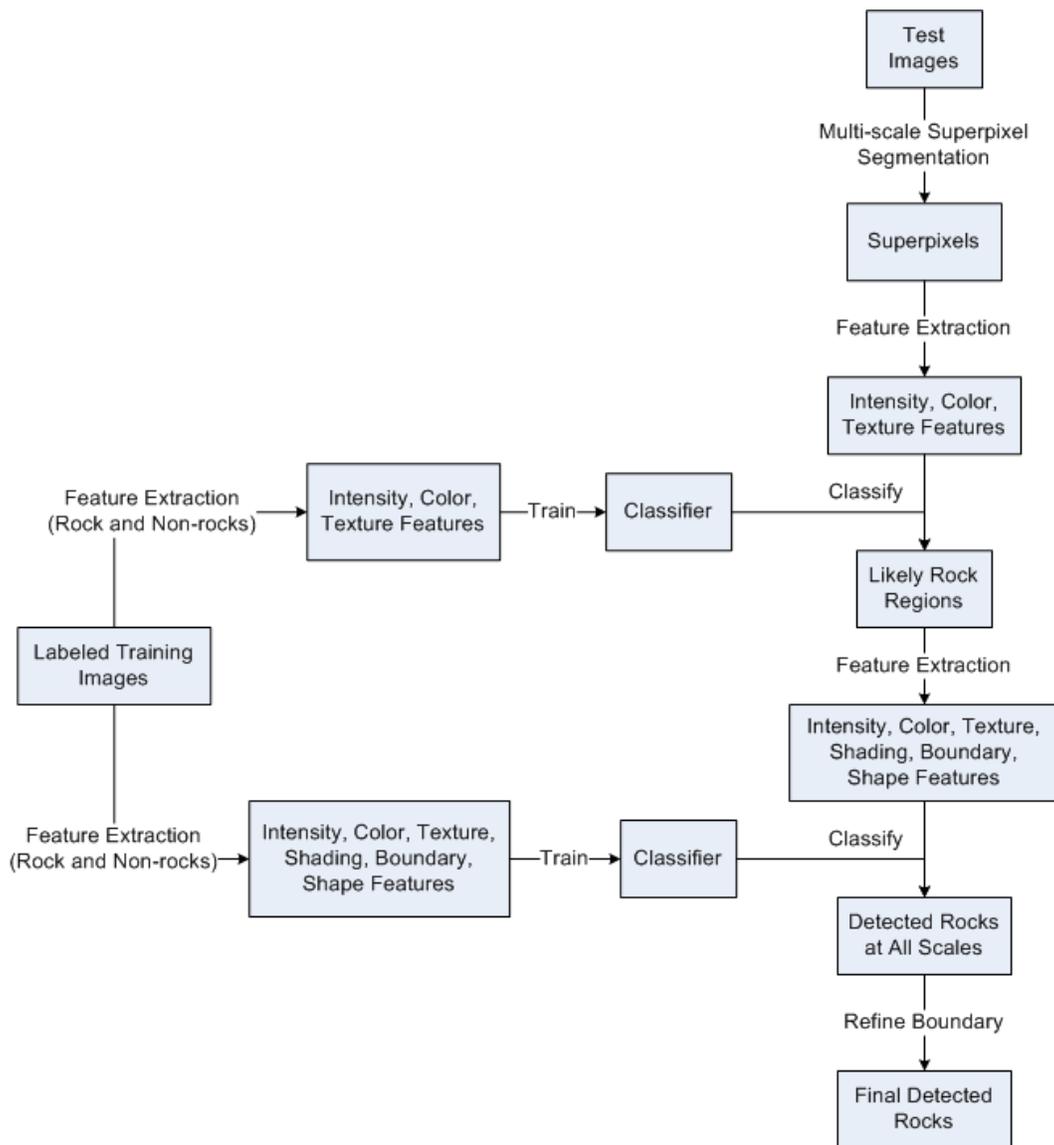


Figure 26: Overview of the rock detection and segmentation algorithm.

first test are examined as possible rocks using the second classifier with intensity, color, texture, shading, boundary and shape features. Collisions between rocks detected at multiple scales are resolved. This process is outlined in Figure 26. This section presents the details of this algorithm.

### 3.2.1 Superpixel Segmentation

As we are attempting to find numerous rocks in each image and cannot expect to find a perfect segmentation algorithm that places each rock in its own segment, a superpixel segmentation is an appropriate strategy. Each of the regions produced from a segmentation algorithm is known as a superpixel. On one extreme, one could treat each image pixel as its own superpixel. With this strategy, it would be certain that some combination of superpixels would accurately cover each rock

region in the scene. However, a more computationally feasible approach with larger superpixels can still retrieve rock boundaries while dramatically reducing the space of superpixel configurations to consider.

A few different segmentation approaches were tried to find the most appropriate for our work. Mean shift [10], Blobworld [4], efficient graph-based [20], pairwise affinities [21] and a variation of normalized cuts [41, 40] segmentations were each evaluated. Blobworld operates by forming a feature vector for each pixel. The features are representations of color and texture in the region surrounding the pixel. These feature vectors are then clustered using the Expectation-Maximization algorithm. Each resulting cluster is a segment. Mean shift segmentation is used in the EDISON system [9]. This two-step technique consists of a mean-shift filtering of the image followed by clustering of the filtered data points. Efficient graph-based segmentation operates on the data points in feature space and uses adaptive thresholding while performing single linkage clustering. Fowlkes' pairwise affinities approach to segmentation combines region and boundary cues. The region cues consist of the similarity in brightness, color and texture between patches. The boundary cue measures the presence of a contour between image pixels. An optimal affinity between pairs of pixels is learned using these cues. Mori's algorithm uses normalized cuts [48] in combination with a boundary detector [38] to provide a contour grouping cue.

Results from these five algorithms are shown for a simple and a more complicated rock scene in Figures 27 and 28. Parameters for each algorithm were experimented with to obtain decent results. Although mean shift does a good job at locating the boundaries around the edges of rocks, it divides each rock into many different segments, even with the minimum region size set fairly high. Blobworld does a poor job at localizing the boundaries of rocks. Although it outlines each rock, this outline is often not located on the actual rock boundary. It also creates a lot of spurious segments surrounding small uneven regions in the background sediment. Efficient graph-based segmentation, for the most part, does a better job than mean shift at localizing boundaries, however, a lot of small spurious regions are created. Fowlkes' segmentation algorithm does a very poor job at identifying rock boundaries. This is very evident with 200 superpixels. Even with 1000 superpixels, the rock boundaries still do not lie on region boundaries. Perhaps the algorithm parameters were not tuned properly for this particular task. Mori's normalized cuts method produces the best results of the algorithms tested. Even with only 100 superpixels, most rock boundaries are identified in the simple scene in Figure 27. Most of the rocks are split into only a few superpixels, an important result necessary to make the later parts of the approach taken in this work feasible. In the more complex scene in Figure 28, more superpixels are required to identify most of the rocks. In both scenes, the boundaries appear to be localized most accurately with Mori's algorithm, with each rock divided into the least number of segments.

We begin with a raw color image. The normalized cuts algorithm requires that the user specify the approximate number of superpixels to create. We desire that each rock is divided into only a few superpixels; to achieve this for many sizes of rocks we segment the image at four different scales as shown in Figures 29 and 30. Thus, at one of these scales, each rock should be divided into only a few superpixels. We are required to subdivide each image into sub-images because of memory constraints in the superpixel segmentation algorithm.

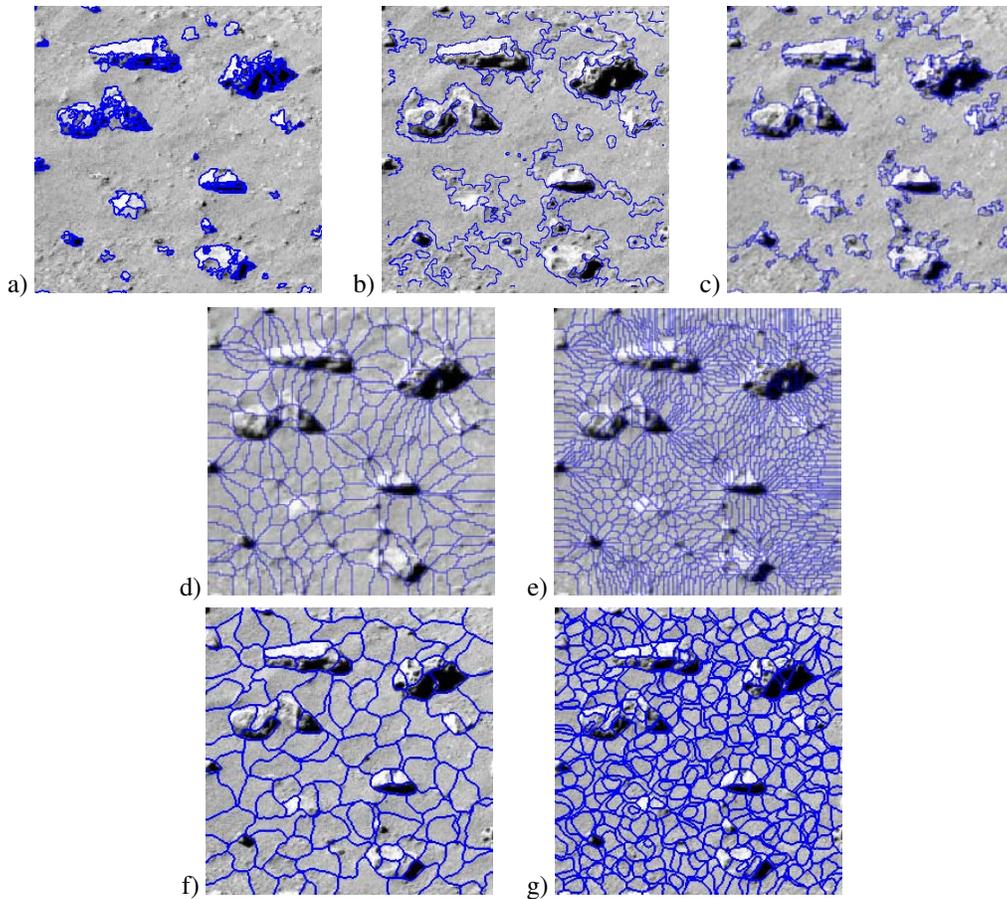


Figure 27: Sample results from a) mean shift, b) blobworld, c) efficient graph-based, d) pairwise affinities with 200 superpixels, e) pairwise affinities with 1000 superpixels, f) normalized cuts with 100 superpixels, g) normalized cuts with 500 superpixels. (Image courtesy of NASA/JPL-Caltech.)

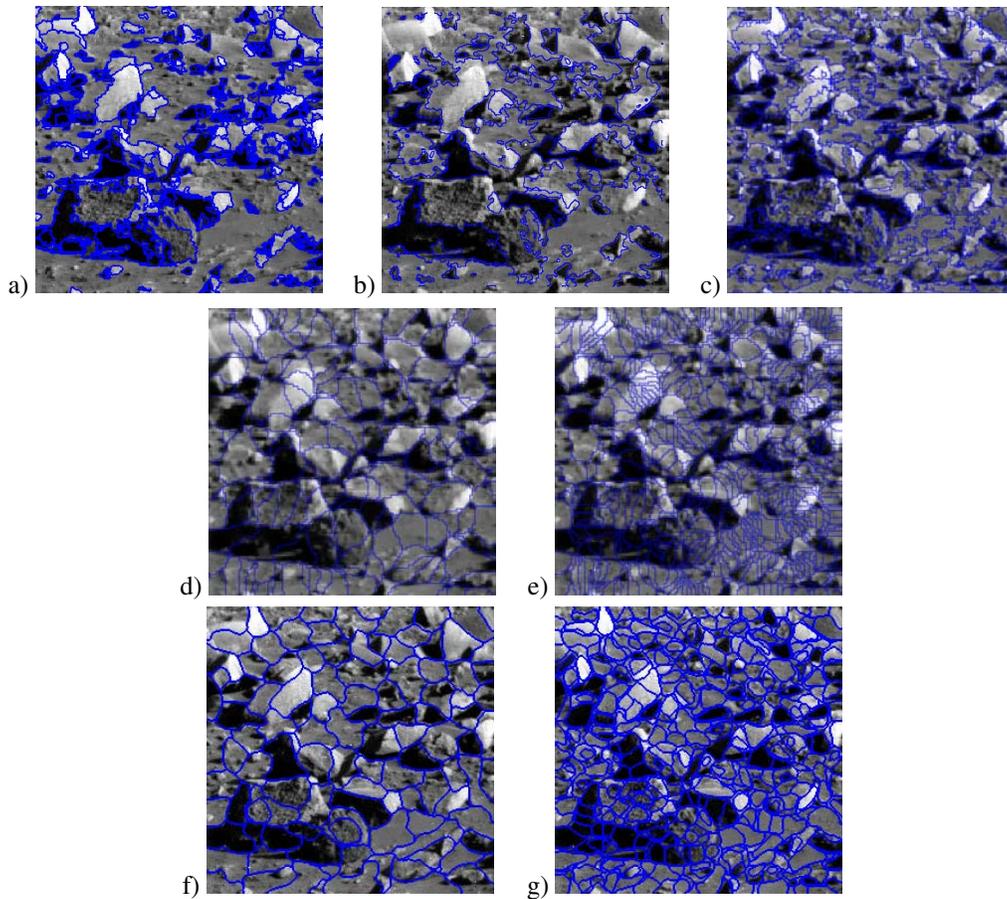


Figure 28: Sample results from a) mean shift, b) blobworld, c) efficient graph-based, d) pairwise affinities with 200 superpixels, e) pairwise affinities with 1000 superpixels, f) normalized cuts with 100 superpixels, g) normalized cuts with 500 superpixels. (Image courtesy of NASA/JPL-Caltech.)

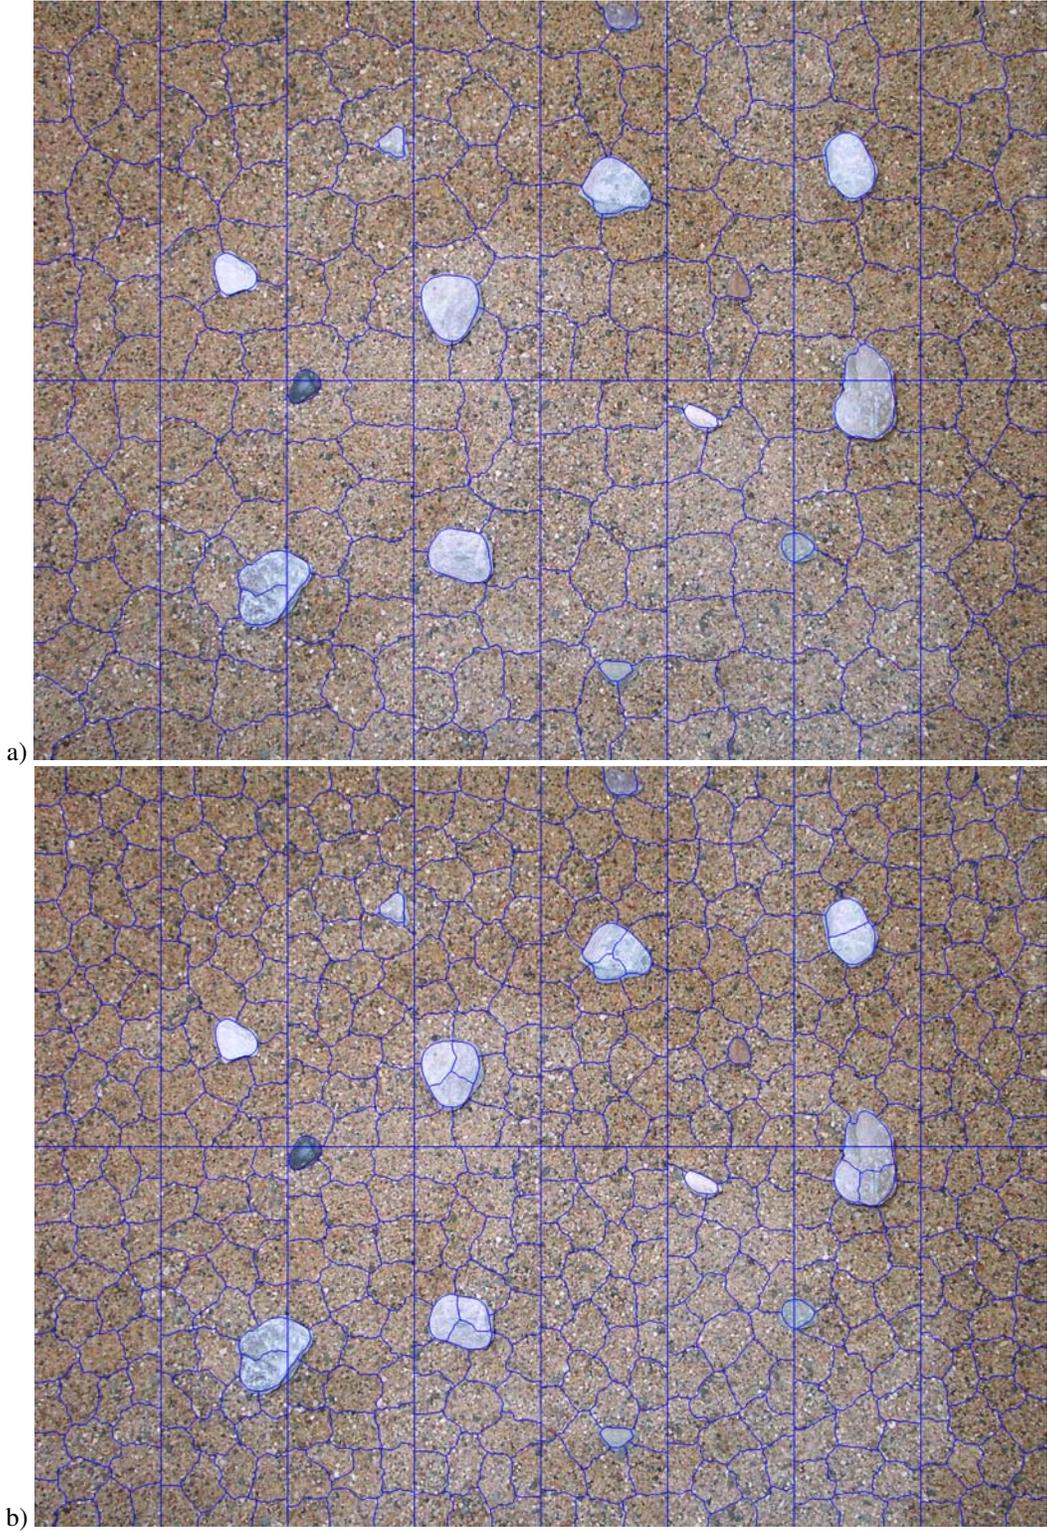


Figure 29: Superpixel segmentations at multiple scales: a) 256 and b) 512 superpixels per image.

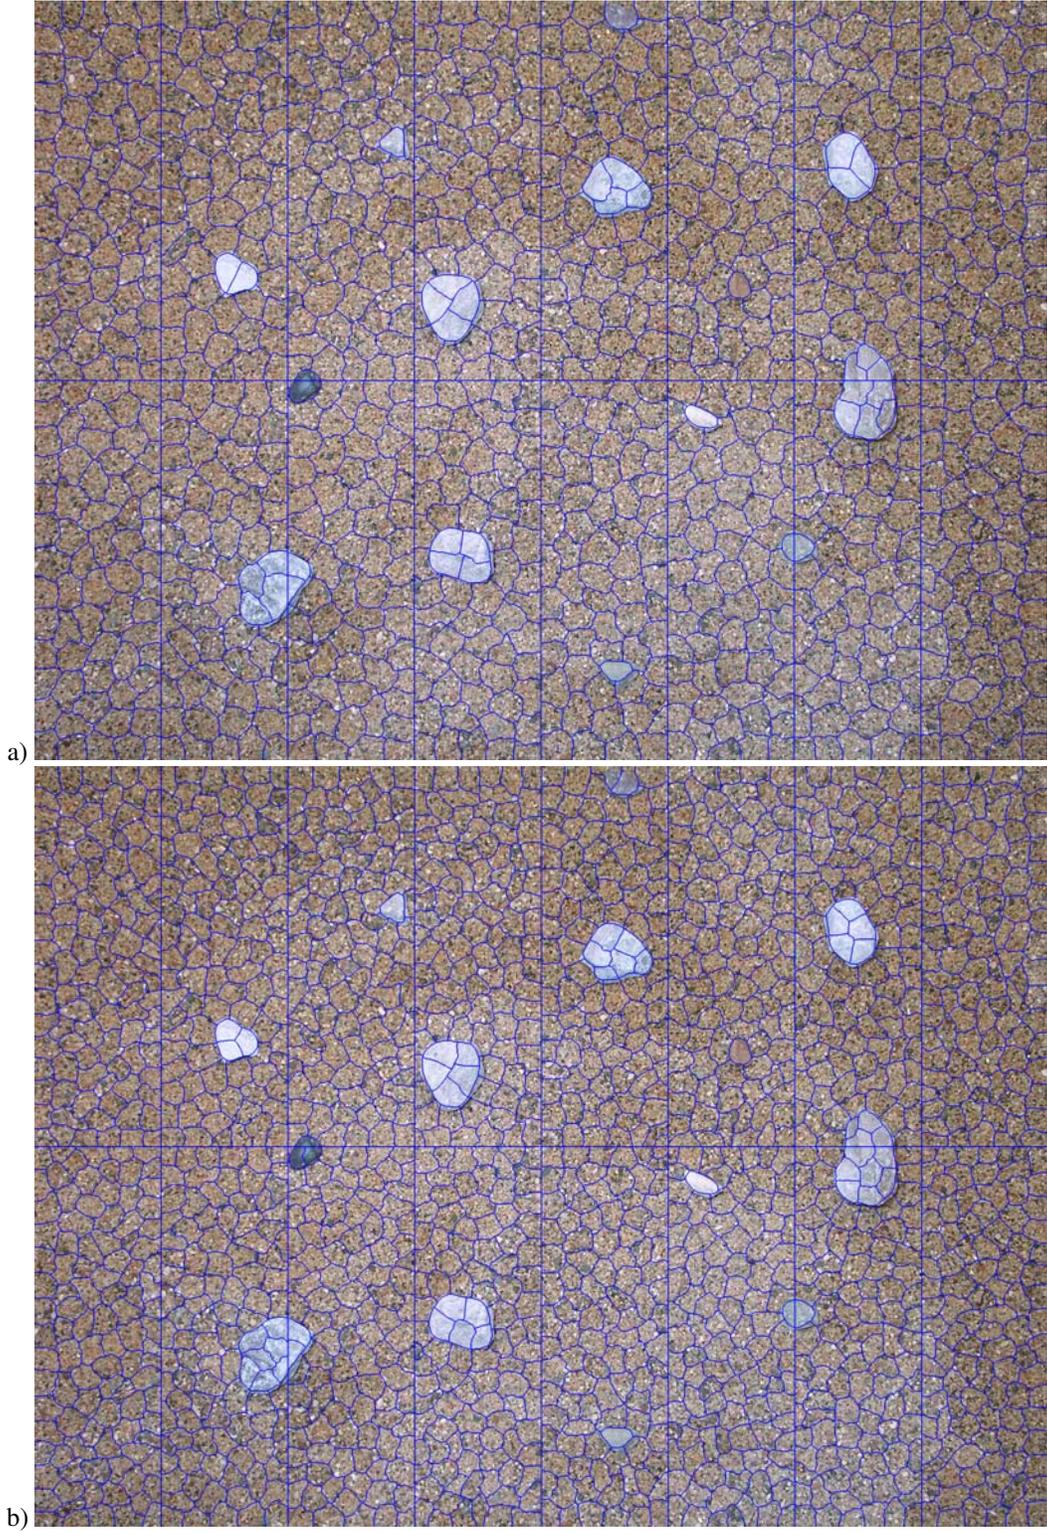


Figure 30: Superpixel segmentations at multiple scales: a) 1024 and b) 2048 superpixels per image.

### 3.2.2 Feature Extraction

We now look at methods for characterizing each superpixel or groups of superpixels. Some of the features used here were developed in Section 2. Others are presented here for the first time as they are only applicable to the detection of rocks, not to their classification.

Some of the region features characterize the region itself, while others measure the difference between the region and its local neighborhood. We define the neighborhood of a region to be its bounding box plus an extra margin of 20 pixels and excluding the region itself. This serves the purpose of distinguishing a rock from the background.

#### 3.2.2.1 Intensity and Color

Measures used for intensity are described in Section 2.1. We use the mean intensity of the region, the variance in intensity of the region, and the absolute value of the difference in mean and variance of intensity between the region and its local neighborhood. We also compute histograms of intensities of the region and its neighborhood. The difference between these two histograms, using either  $\chi^2$  or Euclidean distance, also forms a feature. The same thing is done using the color histograms in RGB, HSV and CIELAB.

#### 3.2.2.2 Texture

In Section 2.2 a few different methods for representing texture were described. Here we select the co-occurrence statistics and textons methods to represent texture. The fractal dimension is not chosen as it is computationally intensive to compute, and the directional histogram method is also discarded as it characterizes a specific property of rocks rather than helping to identify whether a region contains a rock.

Using the co-occurrence statistics method in Section 2.2.2.2, we compute contrast, correlation, energy and homogeneity. These four features are used to characterize the texture of the region.

For the texton method described in Section 2.2.2.4, we use the MR8 filter bank. We use 32 textons and so the texton histogram is a 32-dimensional vector. A texton map is formed for each image, from which the histogram can be computed for any given region. A sample is shown in Figure 31. Notice that many of the rocks are clearly visible in this image, while the background sediment has a fairly uniform texture. Thus, the texton histograms for the rocks should be distinguishable from the background sediment.

Three texture features appear in the final feature vector. The first is the 32-dimensional histogram giving the frequency of occurrence for each texton. The second is a  $\chi^2$  distance measure between the texture histogram of the interior region and the texture histogram of its local neighborhood. The third uses Euclidean distance instead of  $\chi^2$ . These scores provide both a measure of the region texture itself, and the degree to which its texture differs from the background texture.

#### 3.2.2.3 Shape

We extract nine shape measures for each region: elongation, circular variance, elliptic variance, angularity, Diepenbroek roundness, compactness, circularity, convexity perimeter and convexity area. These are each detailed in Section 2.3. These measures are simple to compute, permitting shape information to be evaluated efficiently. From examining the superpixel segmentations in Figures 29, those containing rocks tend to be more round and more convex. The hope is that these shape measures will characterize this tendency.

#### 3.2.2.4 Shading

Rocks are typically convex objects sticking up from a more or less flat bed of background material. For outdoor scenes with natural sunlight, the illumination is very directional, causing highlights and

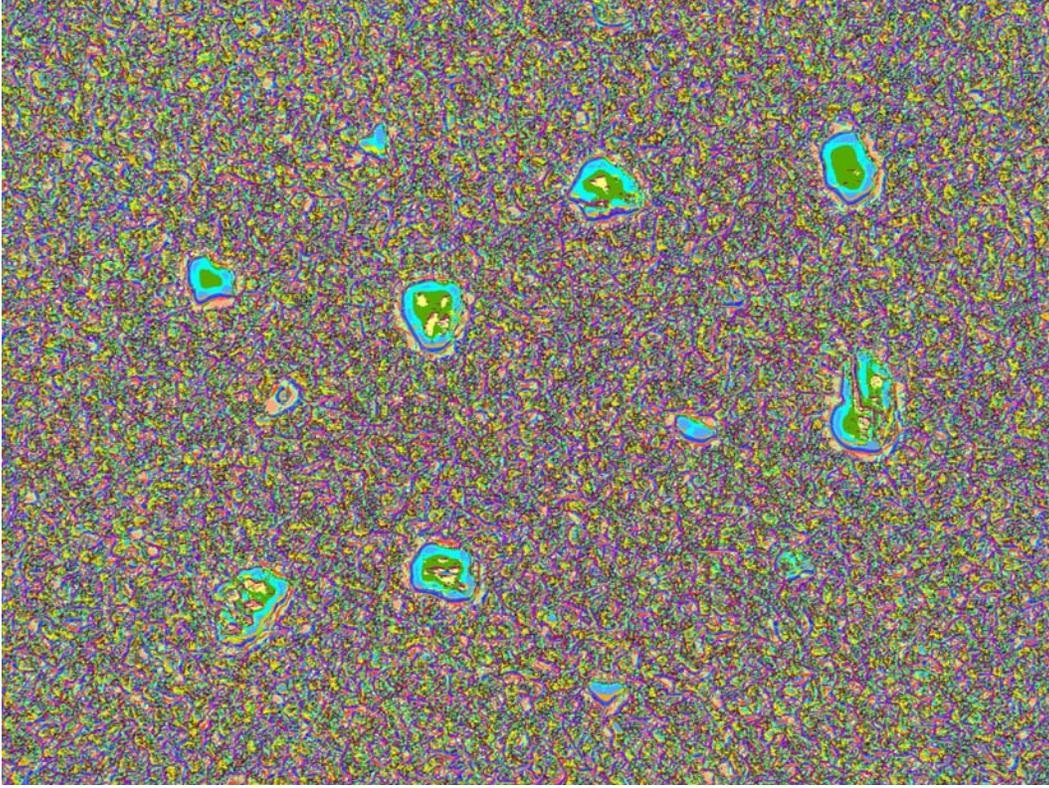


Figure 31: The texton map for the sample image used in Figure 29, computed using the MR8 filter bank, showing the maximally-responding texton at each pixel. Each color represents a different texton.

shadows. The highlight is on one side of the rock, the shadow on the other side and a shading gradient in between. Thus there is an approximately linear relation for pixel intensities of the rock:

$$I(x, y) = ax + by$$

where  $I(x, y)$  is the image, and  $x$  and  $y$  are pixel locations. We first normalize the rock pixel intensities to have zero mean and unit standard deviation, and normalize the pixel locations ( $x$  and  $y$ ) in the same manner). Then we can compute best-fit values for  $a$  and  $b$  with linear regression. For a rock region,  $a$  and  $b$  will have larger (positive or negative) values and for non-rock regions they will be close to zero. The sum squared error to the best linear fit also measures the strength of the gradient fit. The computed  $a$ ,  $b$  and gradient error values all form features for the rock.

In some cases, however, the scene lighting may be less directional or may be directly overhead. In this case, rocks will have a smaller highlight on top and be darker around the edges closer to the ground. To characterize this, we compute the mean pixel intensity of the outer ring of pixels that lie within 7 pixels of the edge. We do the same for the remaining pixels in the interior. By subtracting the mean intensity of the outer region from the mean intensity from the inner region, we expect to obtain a large positive number for rock regions. For non-rock regions, the difference will be closer to zero.

Yet another way to characterize the extent of the shading when the lighting is more diffuse or directional overhead is to fit a quadratic gradient. Assuming that a rock is somewhat spherical, it will be lightest in the center and decreasing in intensity quadratically towards the edges. Thus we fit the

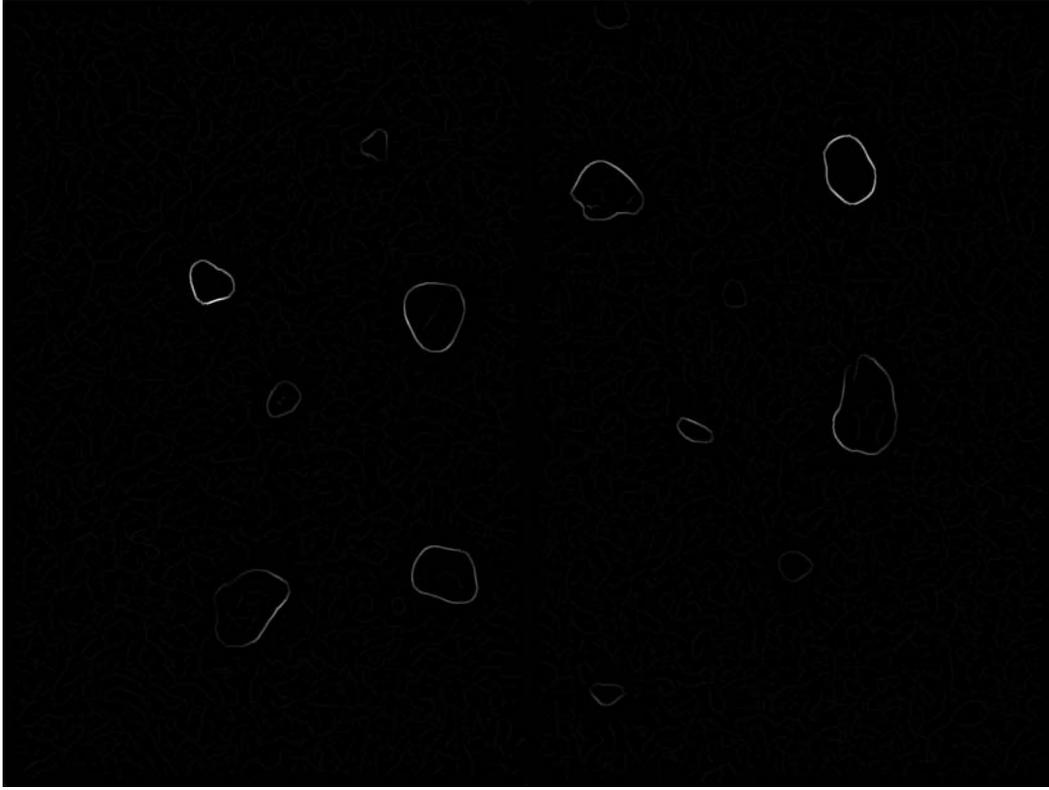


Figure 32: Probability of a boundary for the same image used in Figure 29. Lighter colors correspond to a higher probability.

following quadratic gradient:

$$I(x, y) = ax^2 + by^2 + cxy$$

The sum squared error to this gradient forms another shading feature.

### 3.2.2.5 Boundary Contours

The presence of a visible boundary around the edge of a region can be a powerful cue to the detection of a rock. A boundary exists when there is a transition in intensity or texture from one region to another. We use the approach of Martin, Fowlkes and Malik [38] to natural image boundary detection.

With this method, brightness and texture gradients are computed for each image pixel at multiple orientations. For each pixel, a circle of radius  $r$  is drawn and divided along the diameter at an angle  $\theta$ . The gradient then compares the brightness and texture content in each half-disk. A large difference indicates a boundary in the image along the disk's diameter. This is done at multiple scales for radius  $r$ , and at multiple orientations  $\theta$ . Combining this information, the probability of a boundary existing at each pixel is computed. An example of a boundary probability image is shown in Figure 32.

The contour strength for a given region is then computed by summing the boundary probabilities around the edge of the region up to a 7 pixel margin and normalizing by its perimeter.

### 3.2.3 Rock Detection

At this point we have a multi-scale superpixel segmentation of rock images and a number of features that can be computed on superpixels. The next task is to detect rocks by merging superpixels. We first discard as many background superpixels as possible with the use of a simple classifier. Then the remaining superpixels are examined in numerous combinations of merging to identify the most likely rocks.

#### 3.2.3.1 Training Classifiers

The first step in this process is to create a classifier for rock vs. non-rock regions. A set of training images is used, with the rocks segmented by hand. A random set of superpixels from the multi-scale segmentation of these training images is also selected, such that none of these regions overlap with known rocks. This set of rock and non-rock regions is used to develop two classifiers.

Each of the features described in the previous section is evaluated on this set of regions. This forms a feature vector for each region. Each feature is normalized to zero mean and unit standard deviation. The first classifier we form uses the intensity, color and texture features described in the previous section. A Support Vector Machine (SVM) is trained on this data to form a simple rock vs. non-rock classifier. The second classifier uses all features previously described: intensity, color, texture, shape, shading and boundary contours. Again, an SVM is trained on this data to form a powerful rock vs. non-rock classifier. The distinction between these two classifiers is that the powerful one is applicable only to full rocks, whereas the simple one is also useful in classifying regions that contain only a portion of a rock. For prediction, when the classifier is provided with a new feature vector, the output will include a posterior probability for each class. Thus, not only can we predict whether a region is a rock, we also have a scoring measure for the likelihood that the region is rock.

#### 3.2.3.2 Identifying Candidate Rock Regions

Now we can use the simple classifier to discard background regions. At each scale in the superpixel segmentation, the classifier using intensity, color and texture features is applied to each superpixel. This produces a probability value of that superpixel containing a rock. All superpixels with a greater than 2% probability of being rock regions are selected for further examination. An example of this is shown in Figure 33.

#### 3.2.3.3 Region-Merging

After selecting only the most likely superpixel candidates, most of the image has been discarded. This leaves many fewer combinations of superpixels to examine in detecting rocks. From the remaining superpixels, we form groups that are adjacent. Each group of adjoining superpixels is called a clump.

Before getting into the details of the region-merging algorithm, we make some assumptions to justify the technique used: rocks are relatively well-spaced and each rock is divided into a maximum of ten superpixels at some scale. Well-spaced rocks keep the clumps as small as possible. Some rocks may be close together or even touching, but the majority are well-spaced. Thus, clumps contain relatively few superpixels. For larger rocks, on a lower scale of segmentation, they consist of only a few superpixels. For smaller rocks, the same is true when examining a higher resolution of segmentation. Larger rocks will be made up of numerous superpixels when looking at the higher resolutions. However, they should be detectable at lower resolution segmentations. Thus, we will discard any clumps consisting of more than ten superpixels. For a clump of up to ten superpixels, it is reasonable to evaluate all possible merging combinations, but anything larger is infeasible as the computation grows exponentially.

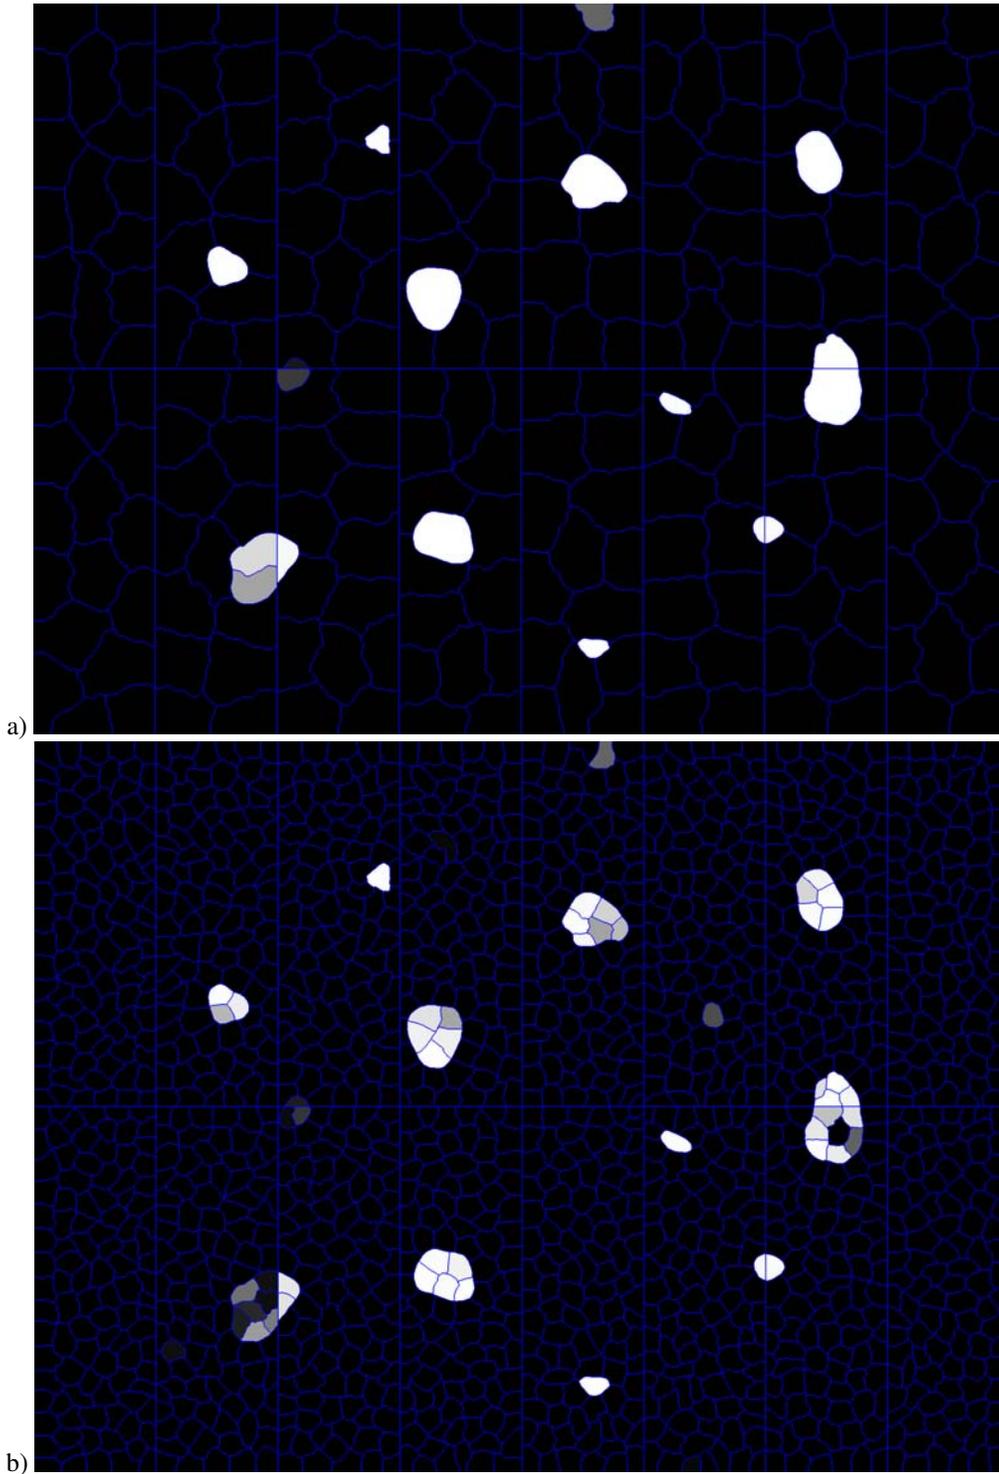


Figure 33: The probability that each superpixel is a rock for the simple classifier using intensity and texture. This is for the superpixel segmentations shown in Figure 29. Lighter colors represent a higher probability for the region. a) 256 superpixels, b) 2048 superpixels.

Working with each segmentation scale at a time, for each clump we examine all possible combinations of superpixels. On each region, the powerful classifier using intensity, color, texture, shape, shading and boundary contours is used, producing a score or posterior probability. For all clumps and all possible region mergings within each clump, the scores are sorted. Starting from the most likely rock, all lower-scoring regions that overlap in part with this region are discarded. Proceeding down the sorted list, this is done until the score falls below a threshold of 20%. All rocks with a score above this threshold that were not discarded due to overlaps are now detected rocks. An example of this is shown in Figure 34.

#### 3.2.3.4 Refine Rock Boundary

Now that we have detected rocks in each superpixel segmentation, the only remaining task is to resolve any overlaps between scales. This is done by selecting the region with the highest score using the powerful classifier. The final results of this process are shown in Figure 35. Typically, when overlapping detections occur at multiple scales, the final detection will contain the rock boundaries that are most likely. If a lower resolution superpixel segmentation does a poor job at locating the boundary for a rock but a higher resolution picks up a better boundary, this better boundary will be used in the final detection.

## 4 Rock Classification

Now that we have looked at different properties of rocks and developed methods for segmenting images and detecting rocks, we can attempt to classify them geologically. The three broad classes of geologic classification are igneous, sedimentary and metamorphic. Igneous and metamorphic rocks originate beneath the earth's surface. Igneous rocks are formed from the cooling and solidification of magma. Metamorphic rocks are formed under high temperatures and pressures. Sedimentary rocks are formed by layer after layer of deposition on the earth's surface at relatively low temperatures and pressures [61]. As long as deposition continues, each layer is under increasing temperature and pressure, converting loose sediment into firmly packed rocks. The parent rock is the main factor in determining the rock's composition, however minerals may be altered or destroyed by weathering and chemical decay [61].

Sedimentary rocks are classified by their physical components. Clastic rocks have grains that were mechanically deposited while non-clastic rocks have grains that were chemically precipitated or recrystallized [61]. The most important potential components of sedimentary rocks are: i) pebbles, sand and coarse silt, ii) mud (clay and fine silt), iii) calcium carbonates and iv)  $\text{SiO}_2$ . In a general classification of sedimentary rocks by grain type, those rich in calcium carbonate are limestone, those with  $\text{SiO}_2$  are chert, mud content is an indication of mudstone and sand an indication of sandstone.

In this section we look at methods for using the albedo, color, texture and shape properties described in Section 2 to geologically classify rocks. The properties form a feature vector for each rock and this vector is used in the classification algorithm.

### 4.1 Related work

Common methods used for supervised learning are k-Nearest Neighbor (k-NN), Decision Trees, Bayes Nets and Support Vector Machines (SVM).

k-NN simply computes the distance between a test point and all points in the training set. Starting at the test point, the first class that has the closest  $k$  data points becomes the assigned class.

Decision Trees usually take the form of binary trees. At each node going down the tree, a condition is tested on a particular feature in the vector. If the test passes, we continue down one branch, if not then down the other. Once a leaf is reached, the final classification is determined.

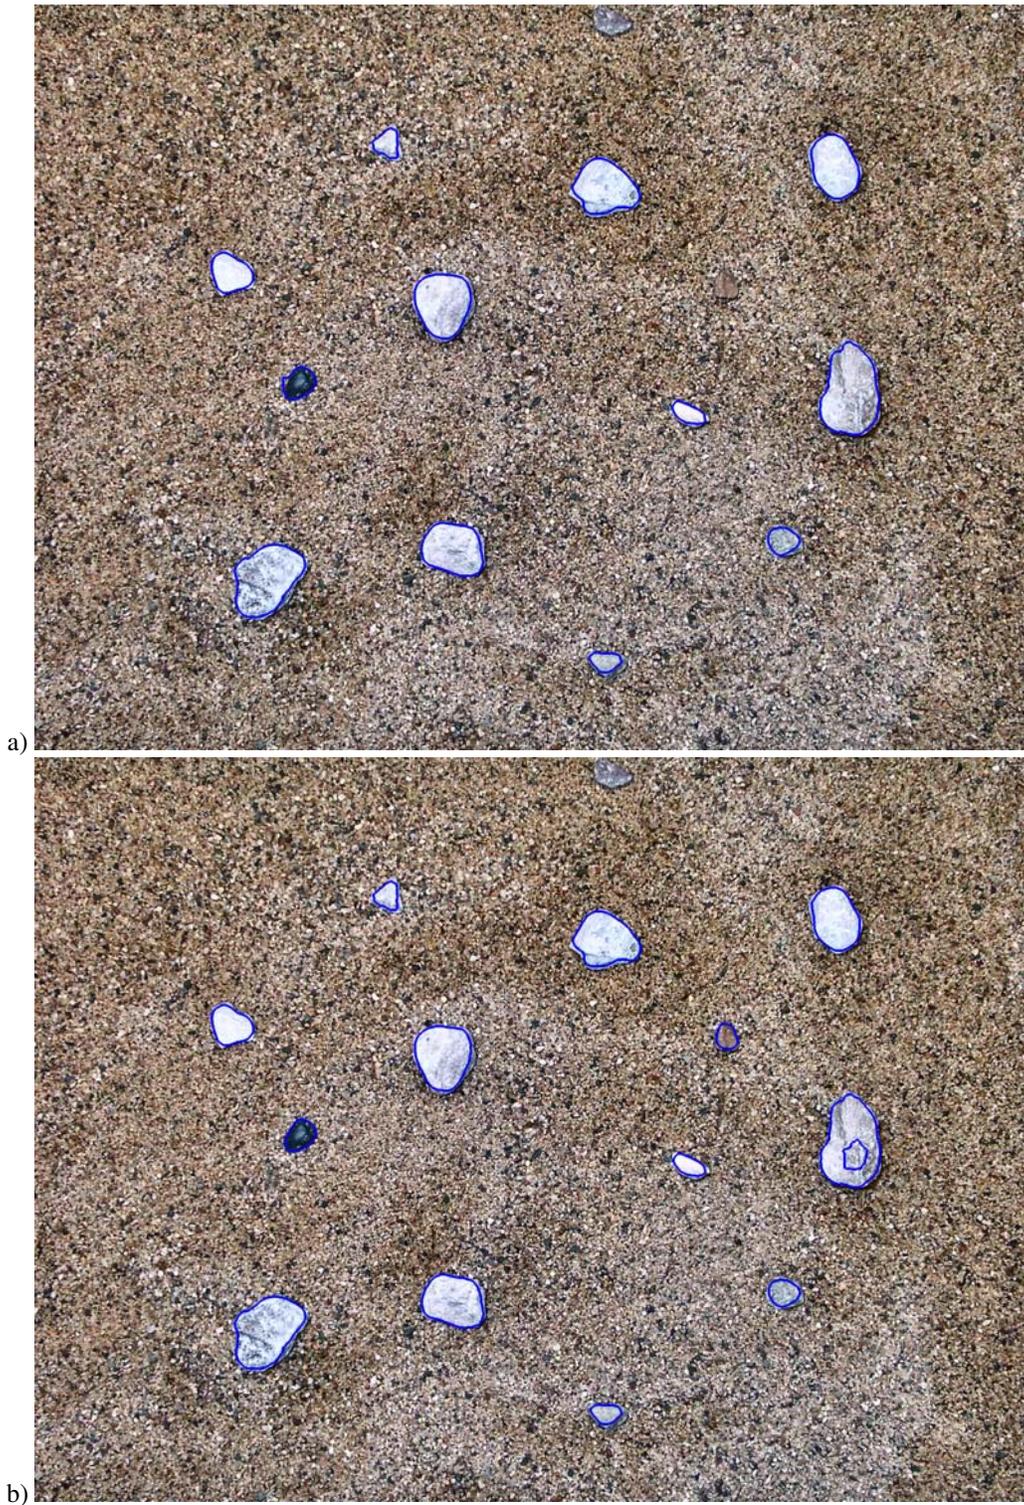


Figure 34: Detected rocks outlines are shown in blue for different segmentation scales: a) 256 superpixels, b) 2048 superpixels.

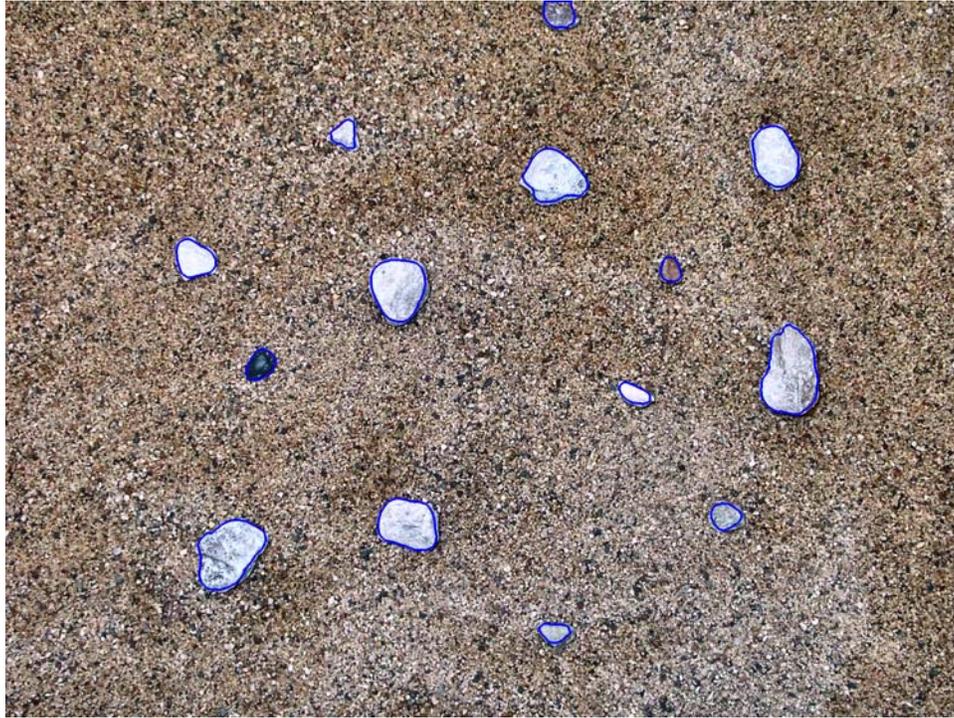


Figure 35: Final rock detections after resolving collisions across the detection scales in Figure 34. Note: the rock at the top of the image was detected by a segmentation scale not shown here.

Bayes Nets represent the dependencies between different features. Each feature in the network has a conditional dependency on some other set of features. These dependencies are represented by conditional probability tables. By knowing the values of some features in the network, others can be predicted, such as the final classification of a feature vector.

A simple SVM with data points belonging to two different classes can be formed by a linear separator. This linear decision boundary will maximize the margin between the decision boundary and the two classes. This simple idea can be adapted to a larger number of classes and into higher dimensions. With a kernel SVM, a set of functions can be applied to the original feature vector in an attempt to find a more accurate decision boundary in a higher dimensional space.

For the specific application of rocks, texture analysis has been used for classification with the use of a k-NN classifier [43, 33, 35, 36].

## 4.2 Approach

In this work we use a couple of simple methods for geologically classifying rocks, as the focus is mainly on the feature sets used, rather than the classifiers.

The first step in developing a classification method is to split the data set into training and testing data. The training data is used in training a classifier, while the testing data is used to assess its accuracy. The training data takes the form of a set of feature vectors, with the features belonging to a subset of those described in Section 2. Each feature vector is a point in a high-dimensional space. The mean and standard deviation for each feature is computed and the data is normalized to zero mean and unit standard deviation.

The first classifier used here is k-Nearest Neighbor (k-NN). Any type of distance metric may be

used with this classifier, but we use Euclidean for simplicity. Different values of  $k$  are tried to find the one that achieves the highest classification accuracy. The algorithm for this classifier is simple. For each test point, the distance to each training point is computed. Starting from closest to the test point and continuing outwards, the classes of the training points are tallied until a class with  $k$  points is reached. This class is assigned to the test point. The k-NN classifier is simple to develop, but can be very slow to compute on large data sets.

We also use a Support Vector Machine (SVM). LIBSVM provides an easy interface to develop a classifier for our purposes [6]. A radial basis function is used for the kernel. The cost and gamma parameters are learned by cross-validation on the training set to find the optimal parameters.

Each of these classifiers provides an easy way to perform geologic classification on rocks from the feature sets computed.

## 5 Experiments

We now examine some experiments performed using the feature extraction, rock detection and classification methods detailed in the previous sections. Using a set of rock images, we will show the effectiveness of the albedo, color, texture and shape feature extraction methods in characterizing rocks. The precision and recall of the rock detection algorithm is examined and the high accuracy in boundary localization proven using a set of rock images, with hand-segmentations as ground truth. We also compare the use of different feature sets to determine the most powerful ones in performing this task. The accuracy in geologic rock classification on a set of rocks classified by an expert geologist is established and different sets of features compared.

### 5.1 Data Set

For this set of experiments, we use a set of rocks previously used in a geology study entitled “The Accuracy of Sediment Size, Shape, and Distribution Measurements from Robotic Geological Images” [58]. Provided from this study are shape measurements of 195 rocks which will be used in determining the accuracy of the automated shape metrics developed in this work.

From this data set, we have eight color images of 2048x1536 resolution, each containing approximately fifteen rocks. We have also been provided with the geologic classifications of 100 rocks from this set and these will be used in establishing the accuracy of the rock classification methods.

The images were taken with a standard digital camera, with all settings on automatic. The rocks have been placed on a bed of sand and each is at a different zoom setting, with individual background sand grains visible in some of the images. A fairly diffuse lighting set-up is used.

Before any experiments are performed, the image is color balanced using Matlab’s adaptive histogram equalization routine. This helps to account for variations in intensity and coloring from image to image that occur as a result of the camera automatically determining the best settings.

### 5.2 Rock Features

Using a set of rocks that were hand-segmented from images, we first look at the use of albedo, color, texture and shape features.

#### 5.2.1 Albedo and Color

As a simple experiment on the albedo and color features, the measures detailed in Section 2.1 were computed on a set of rocks and the resulting feature vectors clustered using k-means. From examining the clusters produced we can get an idea of the effectiveness of these measures.



Figure 36: Four sample classes as a result of clustering the mean intensity feature using k-means. Each row displays example rocks for a particular class.

Figure 36 shows eight example rocks from each of the four clusters produced when using the mean intensity feature. As can be seen, the rock clusters increase in mean intensity from the top row to the bottom.

Figure 37 provides the results from the same experiment but using the variance in intensity feature. The average variance for the cluster increases from the top row to the bottom. The rocks in the second and third rows have a variance in intensity due to changes in the reflectivity of the rock. Row four mostly consists of rocks with dark areas at the edges produced by sharp corners.

Next, Figure 38 performs this experiment on intensity histograms using four clusters. The first row consists of rocks that are uniformly dark. The second row has rocks that are light, but still somewhat darker than those in the fourth row. In the third row, the rocks are somewhat darker, but also have a greater variance in intensity. The fourth row consists of rocks that are uniformly light.

Lastly, Figure 39 shows the clusters produced when using the color histograms from the HSV color space. The first row contains grayish and black rows. Row two has rocks with a red or yellow tint. In the third and fourth rows, the rocks are a very light color, with the third row having a blue tint. The fifth row of rocks are blue. The blue rocks are attributed to the automatic settings on the camera. Ideally, color balancing would have been able to get rid of this; however the effect is still visible.

This experiment provides an indication that the results produced by the albedo and color measures are effective in distinguishing types of rocks.

### 5.2.2 Texture

In a similar manner to the previous section, each texture measure was computed for a set of rocks. The resulting features were clustered into groups. Samples from the clusters are shown in Figures 40 to 43 for the fractal dimension, co-occurrence statistics, directional histograms, and texton histograms using the MR8 filter bank. It is difficult to tell the usefulness of the clusters produced by



Figure 37: Four sample classes as a result of clustering the variance in intensity feature using k-means. Each row displays example rocks for a particular class.



Figure 38: Four sample classes as a result of clustering the intensity histogram feature using k-means. Each row displays example rocks for a particular class. The histogram at the end of the row shows the average histogram for the class.



Figure 39: Five sample classes as a result of clustering the color histogram feature formed using the HSV color space using k-means. Each row displays example rocks for a particular class.

these texture methods from a geology point-of-view, but perhaps a trained geologist would be able to provide insight on this matter.

### 5.2.3 Shape

An algorithm for each of the measures described in Section 2.3 was developed to compute the metrics automatically. Two experiments were performed using these shape metrics. Since sphericity and roundness are established metrics in geology, the first is to compare the results for these with those determined by geologists. The second is to analyze the potential of the other shape metrics for characterizing rocks.

The ideal experiment for sphericity and roundness would be to compare the results of the algorithms developed to values calculated by examining a rock itself and making the measurements necessary to compute the metrics using the definitions. However, this process is very tedious. The data set available includes the determination of these metrics by geologists in a simpler but highly subjective manner, using Crofts' visual chart as shown in Figure 13. This data set is from a geology study on determining the effect of looking at an image of a rock versus the physical rock on a geologist's determination of rock size and shape [58]. The study found that when examining an image of a rock, on average, geologists find the rock to be more spherical and more round than when examining the physical rock. Images of 195 rocks were provided, along with the sphericity and roundness of each, as determined by six different geologists. The observations of the geologists were averaged for each rock. In these experiments we use this data set as well as the images of the rocks in Crofts' chart (Figure 13).



Figure 40: Six sample classes as a result of clustering the fractal features using k-means. Each row displays example rocks for a particular class.



Figure 41: Six sample classes as a result of clustering the co-occurrence statistics features using k-means. Each row displays example rocks for a particular class.



Figure 42: Four sample classes as a result of clustering the directionality features using k-means. Each row displays example rocks for a particular class.

### 5.2.3.1 Sphericity and Roundness Accuracy

In this section we examine the correlation between sphericity and roundness measures determined by geologists and our methods for characterizing rock shape. The first data set we use is Crofts' chart in Figure 13. We use the sphericity and roundness values of 1 to 6 as ground truth and compare with form and roundness measures computed using the algorithms in Section 2.3. We also examine the data set with geologists' determination of sphericity and roundness. We compare both with the geologists' measures from looking at an image of the rock and the actual rock. Results from this experiment are shown in Figures 44 to 51.

For the measures circularity, compactness, convexity perimeter and convexity area, it is not clear whether they should correlate best with sphericity or roundness, and so they are compared with both. The measures compactness, circularity, circular variance and elliptic variance were observed to have a non-linear relation with the ground truth sphericity data, and so an appropriate function is first applied before comparing with ground truth. The function chosen is shown in the axis label on the graph.

We assess the accuracy of the computed measures by computing the correlation coefficient between the computed metric and the geologists' data as ground truth. The correlation coefficient measures the quality of the least squares fit of the data.

Figures 47 and 51 summarize the correlation coefficient between ground truth and each of the metrics. For a metric to be robust, it must achieve a high correlation coefficient on all three data sets. For sphericity, convexity perimeter clearly does not correlate. Convexity area and ellipse error correlate somewhat, but to a lesser degree than the other measures. The remaining metrics - Riley sphericity, elongation, compactness, circularity, circular variance and elliptic variance - all have similar correlation measures on the data sets. Elongation and elliptic variance perform exceptionally well on Crofts' chart, but the performance drops off quite a bit for the other two data sets. This is



Figure 43: Six sample classes as a result of clustering the texton histogram features using k-means. Each row displays example rocks for a particular class.

because in Crofts' chart rocks that are less spherical are more elliptical. However, following other definitions of sphericity, such as Riley sphericity, a rock that is triangular or star shaped will also be less spherical even though it may not be elliptical. Circular variance achieves the highest correlation coefficient on all three data sets.

For roundness, a comparison between metrics is not quite as clear. Convexity perimeter, angularity and the divider dimension have quite a high correlation for Crofts' chart, but this does not hold up on the other two data sets. On the other hand, compactness, circularity, Wadell roundness using the largest corners and the box dimension have low correlation for Crofts' chart but higher for the other two data sets. Wadell roundness following either the standard definition or using the strongest corners performs fairly consistently across the data sets; however, using the standard definition achieves the best results. Diepenbroek roundness achieves a consistently high correlation coefficient across the data sets and appears to be the best performing metric.

We previously mentioned that circularity, compactness, convexity perimeter and convexity area do not clearly belong to the form or roundness category. Results show that circularity and compactness correlate best with sphericity. Convexity perimeter and area, on the other hand, do not characterize sphericity or roundness very well. This is because they measure a different type of property of rock shape.

### 5.2.3.2 Other Metrics

The previous section established the accuracy of sphericity and roundness measures. Here we look at properties of rocks that sphericity and roundness do not characterize, and look at other methods

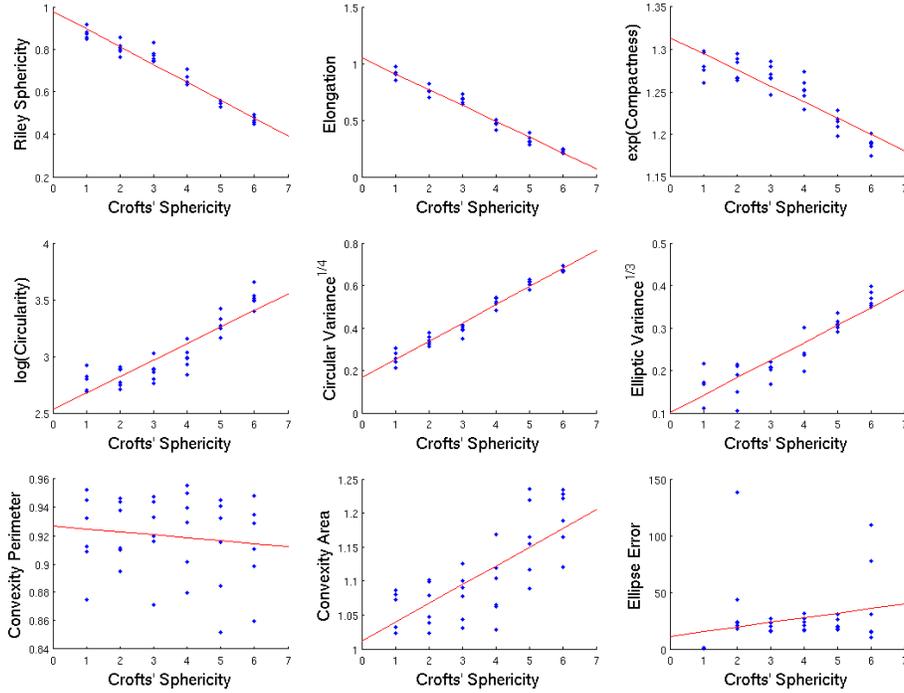


Figure 44: Sphericity for rocks in Crofts chart (Figure 13) vs. computed form measures. The best fit line is shown in red.

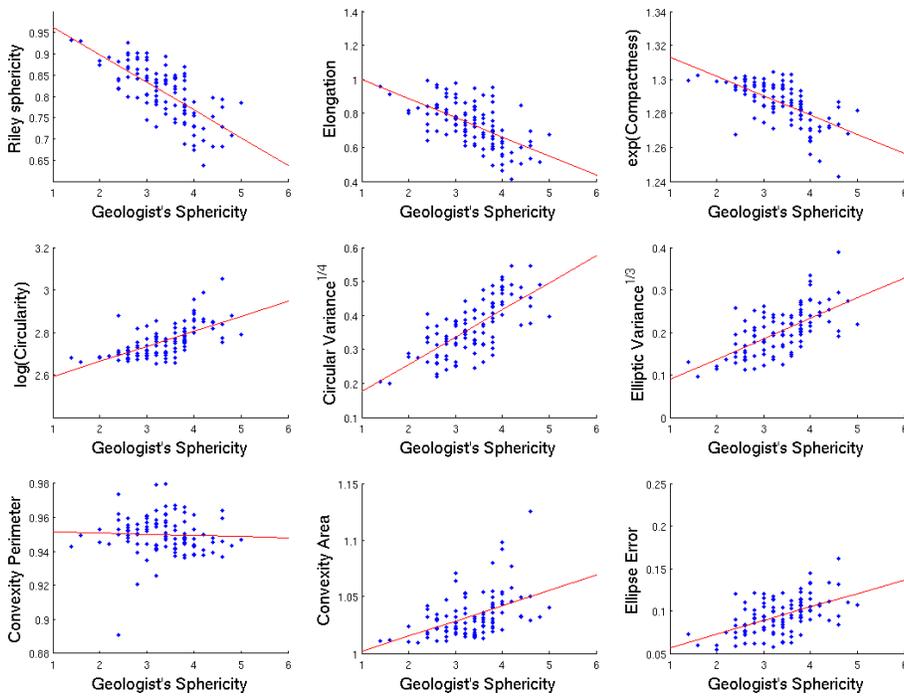


Figure 45: Geologists' sphericity from examining rock images vs. computed form measures. The best fit line is shown in red.

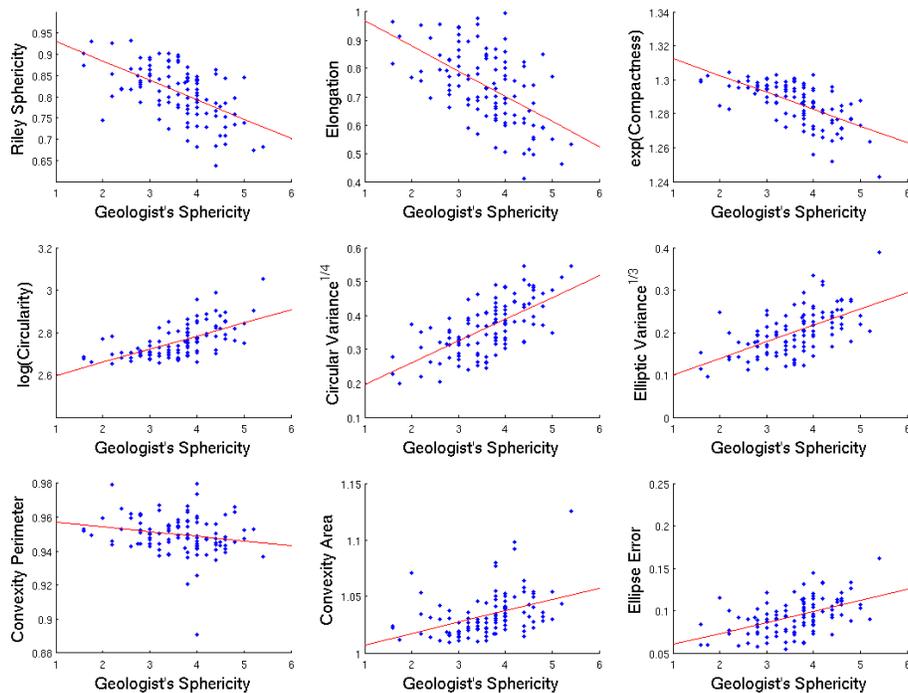


Figure 46: Geologists' sphericity from examining actual rocks vs. computed form measures. The best fit line is shown in red.

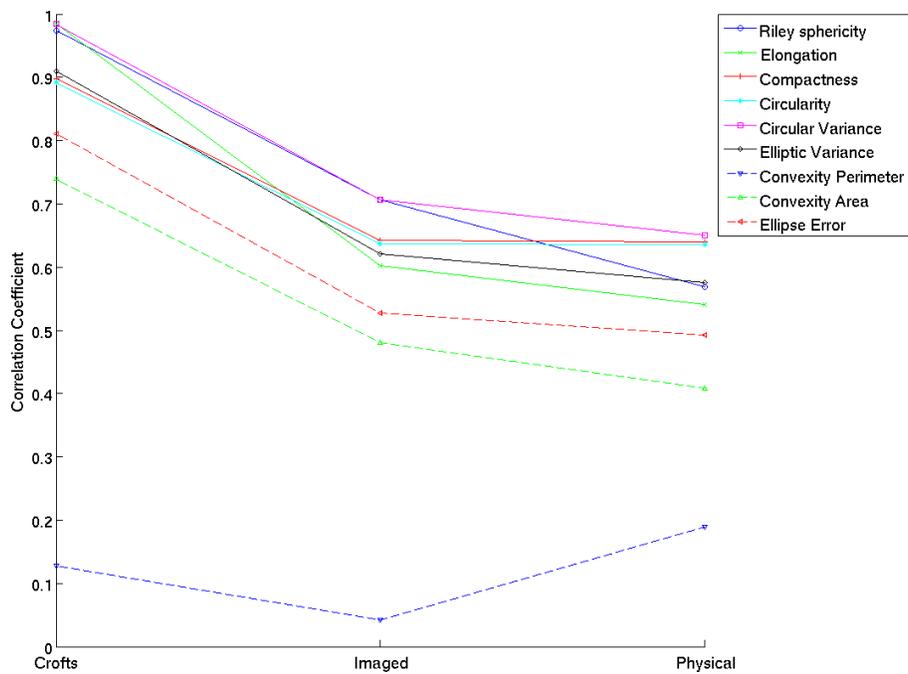


Figure 47: Correlation coefficients between computed form measures and ground truth (data in Figures 44, 45 and 46).

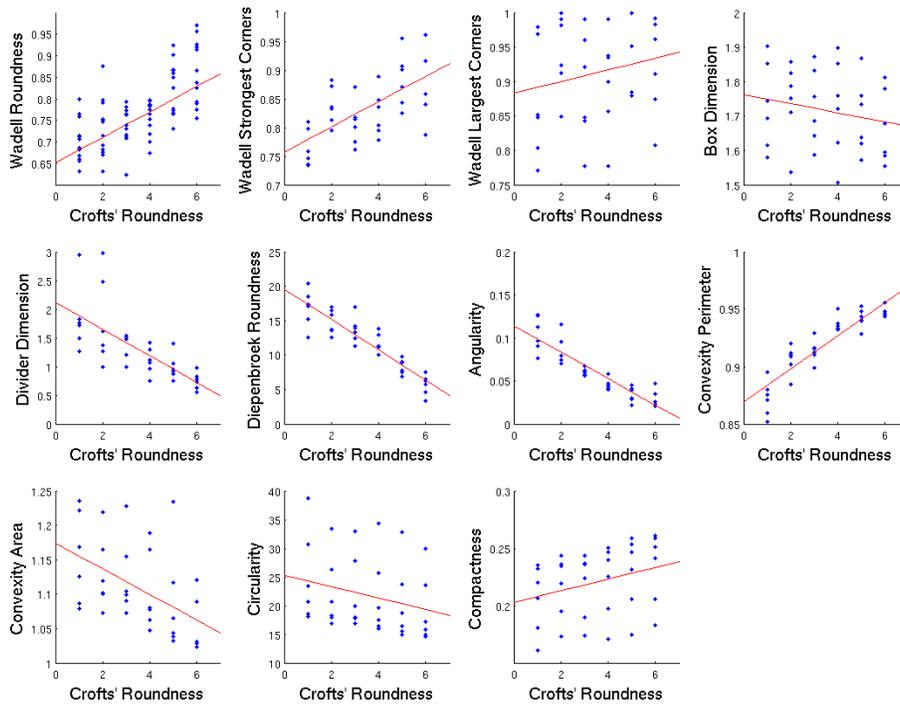


Figure 48: Roundness for rocks in Crofts chart (Figure 13) vs. computed roundness measures. The best fit line is shown in red.

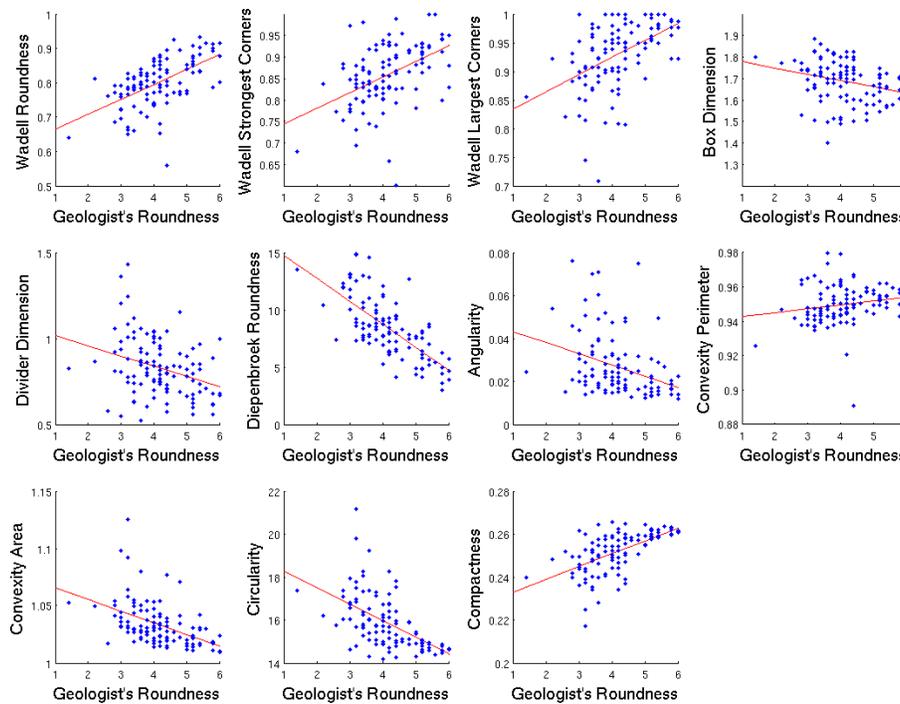


Figure 49: Geologists' roundness from examining rock images vs. computed roundness measures. The best fit line is shown in red.

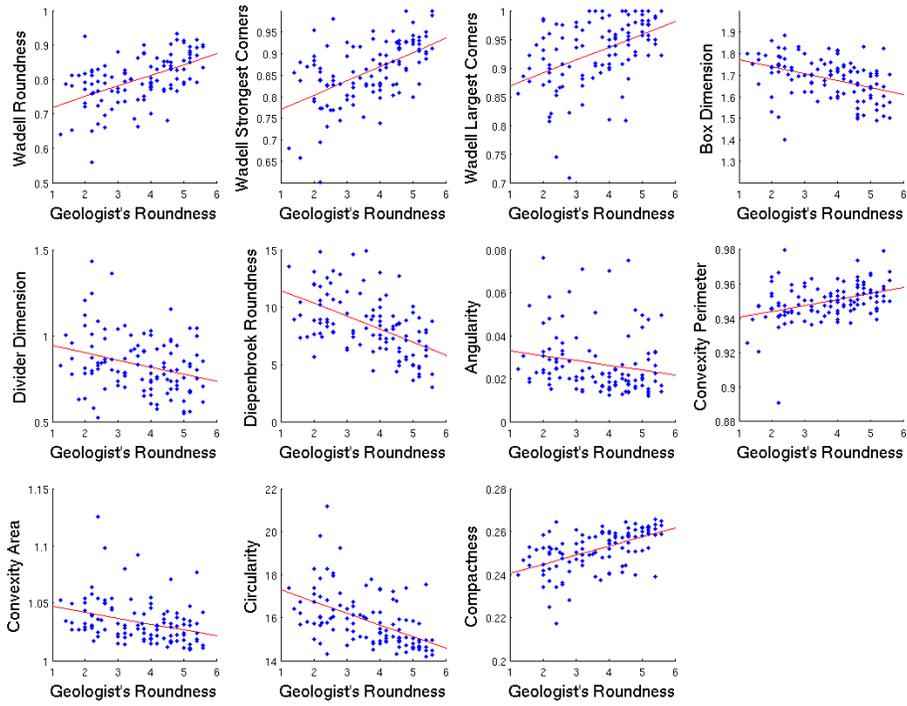


Figure 50: Geologists' roundness from examining actual rocks vs. computed roundness measures. The best fit line is shown in red.

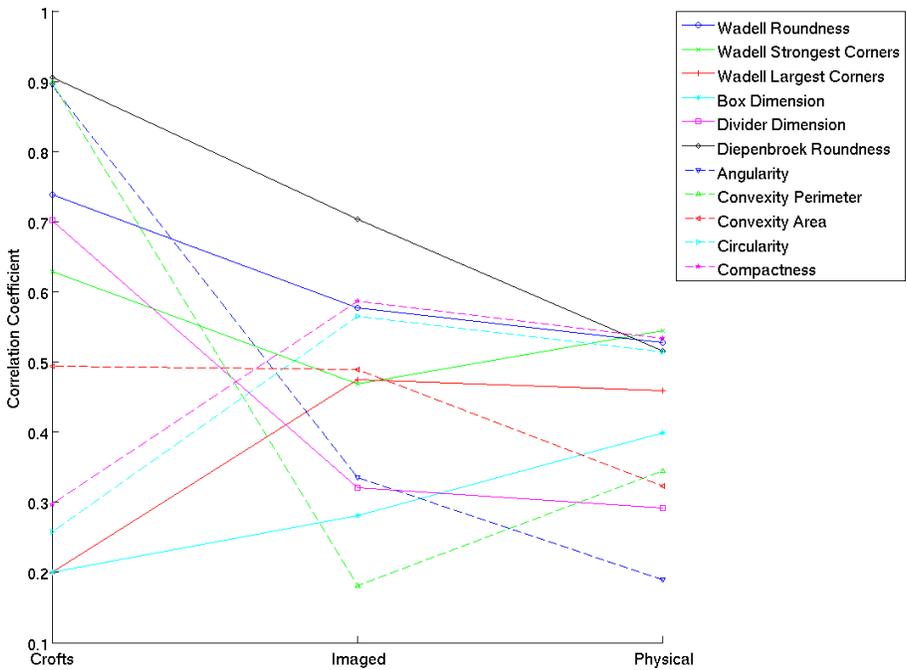


Figure 51: Correlation coefficients between computed roundness measures and ground truth (data in Figures 48, 49 and 50).



Figure 52: Examples of non-convex rocks.

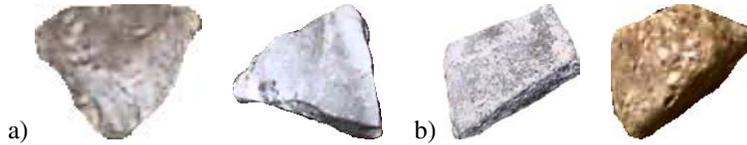


Figure 53: Examples of polygonal rocks: a) triangular rocks, b) quadrilateral rocks.

that might provide useful shape information in addition to sphericity and roundness.

Due to abrasion, parts of a rock can get chipped off resulting in a non-convex surface. This may also happen when a rock first breaks off from its parent outcrop. Thus, convexity may be an important property to measure. The convex perimeter and convex area measures described in Section 2.3 accomplish this task. Figure 52 displays examples of non-convex rocks.

Most of the form measures presented in Figure 2.3 compare the rock boundary to a circle or an ellipse. However, rocks can take the form of other regular shapes. Figure 53 shows examples of non-elliptical regular-shaped rocks. Comparing the boundary of the rock to a triangle, a square or another type of quadrilateral might provide additional useful information about the rocks.

### 5.3 Rock Detection and Segmentation

We now explore a few different methods for assessing the accuracy of the rock detection and segmentation algorithm developed in Section 3. The data set used in this experiment consists of eight images, each containing approximately fifteen rocks. The rocks are placed on a bed of sand. Each image was segmented by hand to obtain a ground truth for the boundary around each rock. An example hand segmentation is shown in Figure 54.

Leave-out-one-image cross validation is used in computing the results. The detectors was trained using seven images and the rocks in each. Non-rock regions were chosen randomly from the superpixel segmentations such that the superpixels do not overlap with a known rocks. Sixty-four superpixels were chosen from each scale in each image, and used as non-rock regions in training.

Sample results are shown in Figure 55. The top image shows very good detection and segmentation accuracy, while the bottom one presents some problems.

The following sections will further examine the accuracy of the detections and boundary localizations using three different methods.

#### 5.3.1 Region Labeling Accuracy

In order to assess the effectiveness of different feature sets in detecting rocks, we simply compute the accuracy of the SVM classifier on different feature sets in distinguishing rock vs. non-rock regions. All eight images are used, with leave-out-one-image cross validation. The rock and non-rock regions used in training the rock detector are used again here. Table 4 provides a detailed listing of the accuracy for all features. The region labeling accuracy is the percentage of regions that are correctly classified as rocks or non-rocks. The full rock detection algorithm is not used here, simply

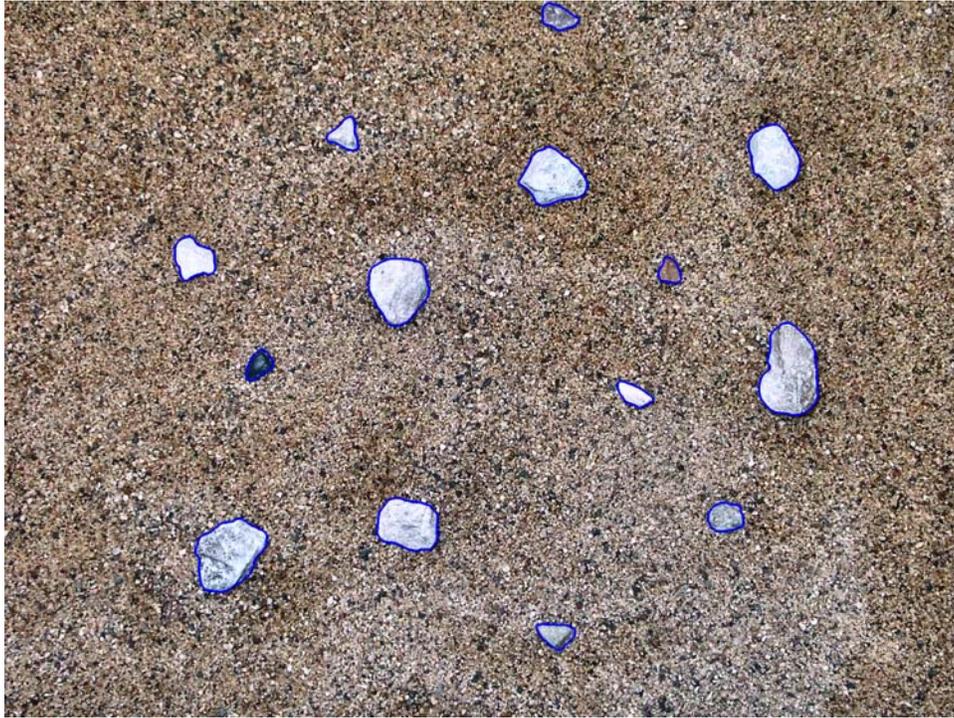


Figure 54: Hand segmentation for a sample image. Each rock is outlined in blue.

a classifier of rock vs. non-rock regions with the purpose of assessing the effectiveness of different feature sets.

The overall region labeling accuracy achieved is 99.6%. Texture proved to be the strongest feature, with the texton approach providing better performance than co-occurrence statistics. The texton histogram out-performed the texton histogram difference methods that compare the region to its neighborhood.

Intensity and color features claimed second place. CIELAB features are the strongest of the three color spaces, likely attributed to the properties of this color space detailed in Section 2.1. The HSV color space also provides stronger features than RGB. Looking at the intensity features, histograms achieved a higher accuracy than mean and variance alone. This is to be expected, as histograms contain much more information.

Shape is the next strongest cue. Angularity and Diepenbroek roundness account for much of this, with the other shape measures having significantly lower region labeling accuracy. This indicates that rock regions can be distinguished from non-rock superpixels by the sharpness of their corners.

The boundary contours cue takes fourth place, while shading is the weakest of the five sets of attributes. The inner/outer mean difference performed the best from this category, while the shading gradient and quadratic shading error helped to boost accuracy of the feature set. The shading gradient would likely be more useful on image sets with directional lighting.

### 5.3.2 Rock Detection Accuracy

The accuracy of rock detection is evaluated with the measures precision and recall. These are computed by first determining whether each detected rock is identified in the labeled data. Correctly detected rocks are defined as those whose detected area overlaps with more than half of the real rock

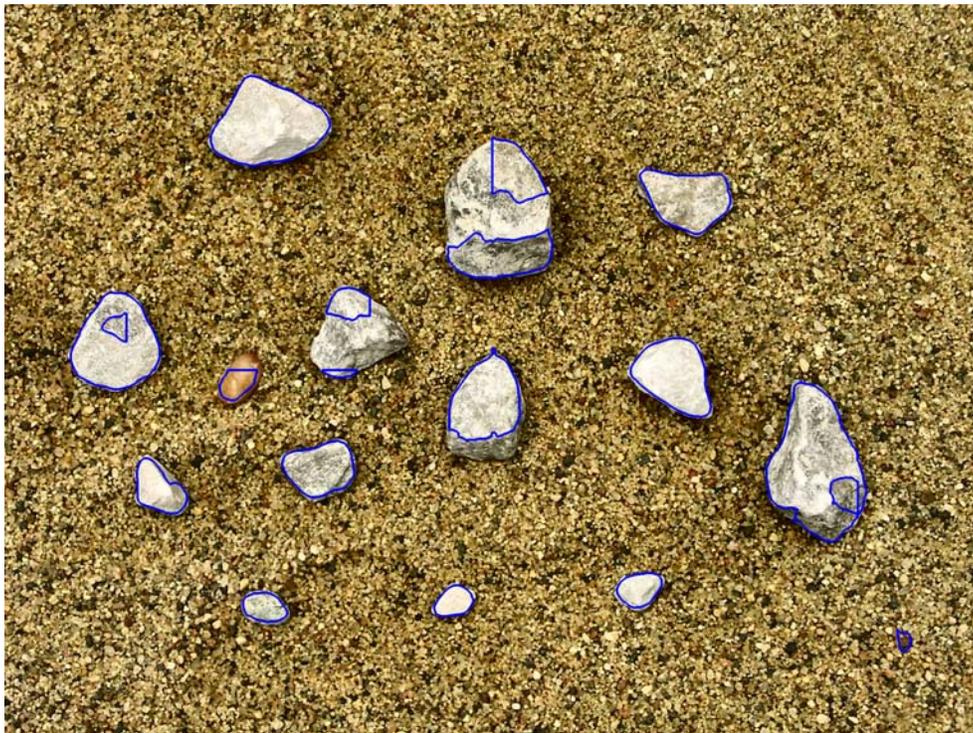
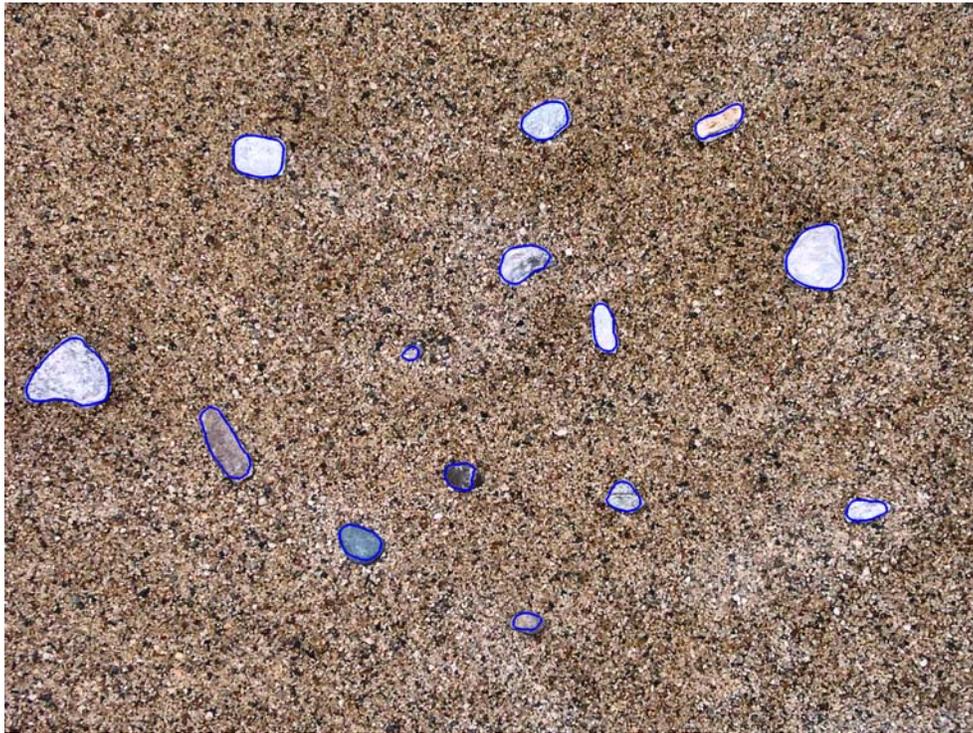


Figure 55: Sample rock detection results. The top image shows good results, while the bottom image achieved poorer detection accuracy.

Table 4: Region labeling accuracy for different features.

<b>Feature Set</b>	<b>Region Labeling Accuracy (%)</b>
<i>All</i>	99.6
<i>Intensity/Color</i>	98.9
Intensity	97.4
Mean	96.7
Variance	94.9
Mean Difference	96.6
Variance Difference	94.3
Histogram	97.4
Histogram Difference $\chi^2$	96.3
Histogram Difference Euclidean	95.3
RGB Color	97.5
HSV Color	98.5
CIELAB Color	98.9
<i>Texture</i>	99.1
Co-occurrence Statistics	98.3
Contrast	98.8
Correlation	95.7
Energy	94.8
Homogeneity	97.3
Textons	99.0
Texton Histogram	99.0
Texton Histogram Difference $\chi^2$	92.7
Texton Histogram Difference Euclidean	94.3
<i>Shape</i>	98.4
Angularity	98.9
Convexity Perimeter	92.8
Convexity Area	92.7
Circularity	92.7
Compactness	93.7
Elongation	92.8
Circular Variance	92.7
Elliptic Variance	92.7
Diepenbroek Roundness	96.5
<i>Shading</i>	97.9
Gradient x-y	94.9
Gradient Error	92.8
Quadratic Error	93.8
Inner/Outer Mean Difference	96.2
<i>Boundary Contours</i>	98.0

Table 5: Precision and recall of the eight images tested.

Image	Precision (%)	Recall (%)
1	100	84.6
2	100	100
3	100	100
4	68.8	73.3
5	80	85.7
6	100	100
7	93.8	93.8
8	45.0	64.3

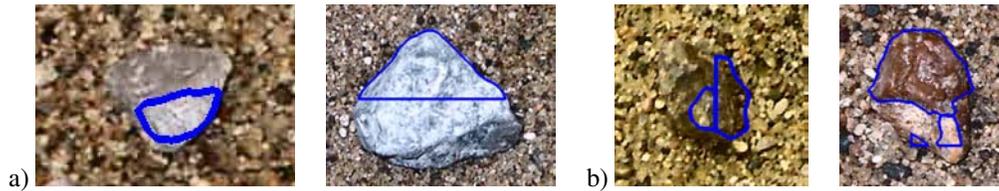


Figure 56: Examples of rocks that were not detected properly. a) Less than half the rock was detected, and so these were not counted as successful detections. b) The rock was split into multiple detections.

area, and whose area is no more than twice the area of the real rock. Precision is then computed as the fraction of detections that are true positives rather than false positives, and recall is the fraction of rocks that are detected rather than missed.

The precision and recall for the eight images is shown in Table 5. The average precision achieved was 85.9% while the average recall was 87.7%. Only one rock in the data set was missed entirely. However, some rocks were split into multiple rocks or only a small portion of the rock was located, decreasing the precision and recall as they may no longer be counted successful detections. Figure 56 shows examples of these cases.

The low precision and recall scores on image 4 resulted from only a small part of three rocks being detected and a fourth rock being split into two unsuccessful detections. Similar situations occurred in other images, however were a more common occurrence in this particular one. For image 8, the magnification is greater than for any of the other images. Thus, in running the rock detection and segmentation algorithm on this image, there was no similarly scaled image in the training set. As measures of texture depend on scale, this is likely an attributing factor for the detection difficulties. This resulted in only small portions of some rocks being detecting, thus not meeting the conditions for a successful detection. In addition, due to the higher magnification, some sand grains in the background were detected as rocks. Any detections smaller than 40 pixels in diameter were discarded, however some were still used in computing the precision and recall scores when no matching ground truth rock had been labeled.

### 5.3.3 Boundary Localization Accuracy

The Chamfer distance measures how closely the detected rock borders match the actual rock borders. This is computed as the average distance from a detected region boundary pixel to the closest boundary pixel on the real rock region. This is done with the assistance of the Matlab function `bwdist` to compute the Euclidean distance transform. Taking a binary mask with ones representing the ground truth rock's boundary, for each pixel the distance transform assigns the distance between that pixel and the nearest non-zero pixel. Now taking the pixel locations for the boundary of the detected rock,

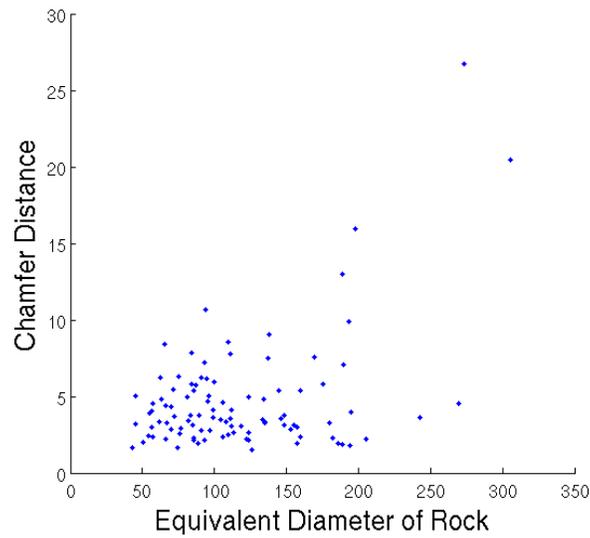


Figure 57: Boundary localization accuracy. The Chamfer distance to the true rock plotted against the equivalent diameter of each rock.

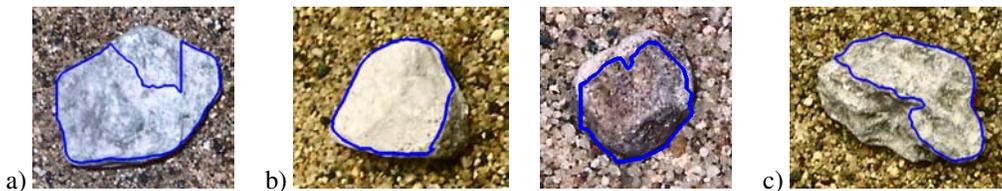


Figure 58: Examples of rocks that were detected, but not properly localized. a) One superpixel was missed in the rock detection. b) The rock intensity is different around a portion of the boundary; this part was missed by the superpixel segmentation algorithm. c) A large chunk of the rock was not detected; this could be due to a set of superpixels being missed or an inaccuracy in the superpixel segmentation.

we sum the value of the previously computed distance transform at these locations and normalize by the perimeter of the boundary.

The boundary localization accuracy for most rocks was found to be quite good. In Figure 57 the chamfer distance to the true rock boundary is plotted against the equivalent diameter of each rock. As can be seen, larger rocks are more likely to have poorer boundary localization. This is usually due to part of the rock not being detected. Figure 58 shows examples of rocks that were detected but not localized properly. Some of these are due to variations in the intensity or texture over the rock.

One factor that might be decreasing the boundary localization accuracy is the quality of the images. As they were taken with a standard digital camera, JPEG compression is used. Thus, the rock boundaries are not as clearly visible than if the images were in a lossless compression format.

From examining the rocks detected at each scale of superpixels, it was noticed that the process of resolving collisions across scales sometimes selects a rock that is more poorly localized than another candidate. No simple method has been found for resolving this as currently the rock with the greatest probability returned from the SVM classifier is selected.

Table 6: Geologic classification accuracy for different classifiers.

Classifier	Classification Accuracy (%)
1-NN	82.4
2-NN	83.4
3-NN	82.9
4-NN	82.1
5-NN	78.9
6-NN	77.6
SVM	86.3

## 5.4 Rock Classification

To assess the performance of the rock classification system developed, we use the same set of images shown in the previous section. The hand-segmented rocks form the training set. Each rock has been classified geologically by an expert geologist into the classes igneous, metamorphic and sedimentary. Sedimentary rocks have also been further broken down into limestone, sandstone, siltstone, quartz and chert, where limestone and chert are chemical rocks and the others are clastic. There are 100 rocks in total, 60 of which are limestone. There are very few igneous and metamorphic rocks in this data set.

In order to form a set with an equal number of rocks in each class and classes as large as possible, the classes chemical and clastic were chosen. There are nineteen rocks available in the clastic class and many more in the chemical class. Thus, nineteen rocks are selected randomly from the chemical class. This set of thirty-eight rocks forms the data set. Leave-out-one-rock cross validation is used in training and testing the classifier. The process of forming the data set and computing the cross validation accuracy is iterated ten times and the results are averaged. Since there are two classes, the simplest classifier of always selecting the same class would achieve 50% accuracy.

Using all features, the classification accuracy for k-NN with different values of  $k$  and for the SVM are shown in Table 6. The SVM achieves the highest accuracy of 86.3%.

For a detailed breakdown of classification accuracy by feature set, we use an SVM. Table 7 shows the accuracy achieved for each feature. Intensity and color prove to be the strongest features, followed by texture.

CIELAB histograms achieve a classification accuracy of 85.3%, the highest of any specific feature. As expected, color features consistently outperform intensity ones. RGB and HSV features achieve comparable results.

For texture, texton histograms using the MR8 filter bank are the strongest feature at 79.5% classification accuracy. Texton histograms with the Gabor filter bank and directionality histograms also provide some use. The fractal dimension and co-occurrence statistics achieve poor results.

Wadell and Diepenbroek roundness both achieve some accuracy in classification. All other shape metrics appear to provide little use. Perhaps with different classes, such as including igneous and metamorphic rocks, the result would be different.

An attempt was also made to perform classification with igneous, limestone and sandstone as classes. However, as only ten igneous rocks are available in the data set, there were insufficient examples to obtain conclusive results.

## 6 Conclusions

In this report we have developed and compared methods for the detection and geologic classification of rocks.

Table 7: Geologic classification accuracy for feature sets using an SVM classifier.

<b>Feature</b>	<b>Classification Accuracy (%)</b>
<i>All</i>	86.3
<i>Intensity/Color</i>	86.1
Intensity Mean and Variance	60.3
Intensity Histogram	56.1
RGB Mean and Variance	73.9
RGB Histogram	74.2
HSV Mean and Variance	67.9
HSV Histogram	73.4
CIELAB Mean and Variance	84.2
CIELAB Histogram	85.3
<i>Texture</i>	76.3
Fractal Dimension	58.9
Co-occurrence Statistics	58.4
Directionality Histogram	61.8
Texton Histogram (Gabor)	67.4
Texton Histogram (MR8)	79.5
<i>Shape</i>	70.0
Riley Sphericity	41.8
Elongation	41.6
Ellipse Error	62.1
Circular Variance	36.3
Elliptic Variance	54.2
Wadell Roundness	75.5
Wadell Roundness (Strongest Corners)	78.2
Wadell Roundness (Largest Corners)	72.6
Angularity	48.7
Fractal Dimension (Box)	56.1
Fractal Dimension (Divider)	38.2
Diepenbroek Roundness	67.6
Compactness	50.8
Circularity	41.6
Convexity Perimeter	35.2
Convexity Area	44.7

## 6.1 Summary

Here we provide a summary of the methods and results of this work.

**Feature extraction** Methods for characterizing rock albedo and color using mean, variance and histograms were explored. The color spaces RGB, HSV and CIELAB can be used in color analysis. These methods were shown to produce useful results when clustering rocks. For the purpose of rock detection and classification, color histograms using the CIELAB color space proved most useful.

For rock texture analysis, methods using the fractal dimension, co-occurrence statistics, directional histograms and textons were developed. The texton approach with the MR8 filter bank outperformed all other texture methods in rock detection and classification.

Numerous shape metrics were described for characterizing rock form and roundness. Circular variance was shown to correlate best with sphericity measures by geologists. Diepenbroek roundness correlated best with geologists' roundness measures. Metrics to characterize the shape of rocks in ways other than the standard sphericity and roundness were also explored. For rock detection, the measures angularity and Diepenbroek roundness performed best. For classification, Wadell and Diepenbroek roundness provided decent results.

**Rock detection and segmentation** A combination detection-segmentation algorithm was developed for locating and accurately defining the boundary of rocks. This method performed very well on the data set, detecting rocks with a high precision and recall. The boundary of most rocks was also accurately localized. However, only small portions of some rocks were detected, decreasing the precision and recall. Smaller rocks were most accurately localized. This method operated on a data set of images containing well-spaced rocks under diffuse lighting.

**Rock classification** Geologic classification of rocks was attempted using a k-NN and SVM classifier. Leave-out-one-rock cross-validation was used in computing the classification accuracy. The SVM performed the best, with color features as the strongest indicator. CIELAB histograms accounted for most of the accuracy in classification. Texton histograms with the MR8 filter bank also achieved good results. The classes chemical and clastic rocks were used in order to provide a sufficient number of training examples.

## 6.2 Contributions

The techniques developed in this report have been successfully used in automated geologic analysis. The important contributions of this work include:

**Powerful object characterization features** Many of the rock features developed here are quite powerful in characterizing rocks. Albedo, color, texture and shape are all important visual measures of objects in an image and important geologic properties of rocks. This investigation into their use in characterizing rocks can also apply to characterizing, detecting and classifying other natural objects and terrain features.

**Automated sphericity and roundness measures for geologists** The circular variance and Diepenbroek roundness metrics developed correlate very well with geologists' measures of sphericity and roundness. As these measures are used so often in geology, a repeatable and reliable automated method for computing them will prove useful.

**Multi-scale rock detection with accurate boundary localization** The rock detection and segmentation routine developed not only locates rocks in an image, but also determines an accurate boundary. Most current rock detection algorithms, and object detection algorithms in general, only place a bounding rectangle around detected objects. The method developed here is necessary for further geologic analysis of rocks. This method also achieves high performance in detecting rocks at multiple scales simultaneously.

**Geologic classification** The geologic classification method described here uses a wider selection of features than most others previously published. The use of these and other features can perform automated geologic analysis of rocks.

### 6.3 Future Work

Although much has been accomplished in this project, there are still numerous directions for future work. Some possibilities include:

**Color balancing** One problem encountered with color resulted from the automatic settings on the digital camera. A better method for color balancing or calibrating images would likely improve the usefulness of the color measures.

**Texton methods and filter banks** A more complete investigation into the use of textons to analyze rock texture would be useful. Numerous different filter banks, as well as the MRF approach, have been used in texture analysis and perhaps some are better than others when applied to rocks. The MR8 filter bank discards potentially useful data when it takes the maximum across all orientations. Perhaps retaining the response from each orientation would improve results.

**Subjectivity of geologists measures** The difficult part of the rock shape analysis performed here is that geologist's measurements of sphericity and roundness are highly subjective, so there is no real ground truth data set to compare to. An interesting experiment would be to measure the sphericity and roundness of rocks by hand and compare those values to the results of the algorithms described here, operating on images of the rocks.

**Effect of viewpoint and camera on shape metrics** The effect of performing these measurements from a single image of a rock should be investigated. Maybe taking images from multiple views and combining the results would produce a more accurate value. The effects of camera resolution and the size of the rock are also important.

**Fourier analysis of shape** The shape metrics described here are just a few of those available. Fourier analysis can be used in ways to describe form as well as roundness. Other methods measuring the curvature of the rock can be used to characterize roundness. These possibilities could also be explored as more alternatives for analyzing rock shape.

**Superpixel segmentation algorithms** The superpixel segmentation algorithm used here was selected after some initial experiments. A more thorough investigation into different methods and the detection results achieved from each is advisable.

**More difficult rock detection data set** Because such high precision and recall are achieved in rock detection, this is an indication that a more difficult data set is needed. Numerous other complications exist which do not occur in the current set of images. Rocks can be overlapping, covered in dust, partially covered in sand or under highly directional lighting causing cast shadows. A data set exhibiting these possibilities would give a better indication of the performance of this algorithm in real-world scenarios.

**Effect of segmentation errors on feature extraction and classification** Section 5.3 shows results for the accuracy in rock detection and segmentation. However, the segmentation for some rocks does not result in an accurately localized boundary. In these cases, the albedo, color and texture features may not be affected too much, however the shape metrics will be. The effect of segmentation errors on these features and the resulting classification has not yet been determined.

**Larger and more diverse rock selection** Although a greater classification accuracy might be desirable, this work provides an indication of the types of features that will be successful in geologic classification. A larger and more diverse data set is desired to further test these methods.

## 6.4 Closing Remarks

The algorithms developed here have been shown to perform the desired tasks well. They are applicable for use in performing geologic analysis of images returned from remote exploration vehicles, such as the Mars Exploration Rovers. With further development of these techniques and improved computational efficiency, they can be used in on-board autonomous science systems. Perhaps a future generation of this technology will some day be in use on Martian rovers.

## References

- [1] P.J. Barrett. The shape of rock particles, a critical review. *Sedimentology*, 27:291-303, 1980.
- [2] H. Blatt. *Sedimentary Petrology*. W. H. Freeman and Company, 1982.
- [3] E.T. Bowman, K. Soga, and W. Drummond. Particle shape characterization using fourier descriptor analysis. *Geotechnique*, 51(6):545–554, 2001.
- [4] C. Carson, S. Belongie, H. Greenspan, and J. Malik. Blobworld: Image segmentation using expectation-maximization and its application to image querying. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(8):1026–1038, 2002.
- [5] R. Castaño, R.C. Anderson, T. Estlin, D. DeCoste, F. Fisher, D. Gains, D. Mazzoni, and M. Judd. Rover traverse science for increased mission science return. In *IEEE Aerospace Conference Proceedings*, March 2005.
- [6] Chih-Chung Chang and Chih-Jen Lin. *LIBSVM, a library for support vector machines*, 2001. Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.
- [7] B.B. Chaudhuri and N. Sarkar. Texture segmentation using fractal dimension. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17:7277, 1995.
- [8] R.J. Cheel. Introduction to clastic sedimentology. In *ERSC 2P10 Course Notes*. Brock University, Ontario, 2005.
- [9] C. Christoudias, B. Georgescu, and P. Meer. Synergism in low level vision. In *International Conference on Pattern Recognition*, volume 4, pages 150–155, 2002.
- [10] Dorin Comaniciu and Peter Meer. Mean shift: A robust approach toward feature space analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24:603–619, 2002.
- [11] R.C. Crida and G. de Jager. Rock recognition using feature classification. In *Proceedings of the IEEE South African Symposium on Communications and Signal Processing*, pages 152–157, October 1994.
- [12] R.C. Crida and G. de Jager. Multiscalar rock recognition using active vision. In *Proceedings of the IEEE International Conference on Image Processing*, pages 345–348, September 1996.
- [13] R.S. Crofts. A visual measure of shingle particle form for use in the field. *Journal of Sedimentary Petrology*, 44:931–934, 1974.
- [14] L. da F. Costa and R.M. Cesar. *Shape Analysis and Classification: Theory and Practice*. CRC Press, 2001.
- [15] M. Diepenbroek, A. Bartholoma, and H. Ibbeken. How round is round? a new approach to the topic 'roundness' by fourier grain shape analysis. *Sedimentology*, 39:411–422, 1992.
- [16] G.R. Drevin. Using entropy to determine the roundness of rock particles. In *Proceedings of 5th International Conference on Signal Processing*, volume 2, 2000.
- [17] G.R. Drevin and L. Vincent. Granulometric determination of sedimentary rock particle roundness. In *Proceedings of the International Symposium on Mathematical Morphology*, pages 315–325, April 2002.
- [18] D.J. Durian, H. Bideaud, P. Düringer, A. Schröder, and C.M. Marques. The shape and erosion of pebbles. *Physical Review Letters*, 97, July 2006.
- [19] R. Ehrlich and B. Weinberg. An exact method for characterization of grain shape. *Journal of Sedimentary Petrology*, 40(1):205–212, 1970.
- [20] Pedro F. Felzenszwalb and Daniel P. Huttenlocher. Efficient graph-based image segmentation. *International Journal of Computer Vision*, 59(2), September 2004.
- [21] C. Fowlkes, D. Martin, and J. Malik. Learning affinity functions for image segmentation: Combining patch-based and gradient-based approaches. In *CVPR*, June 2003.
- [22] J. Fox, R. Castaño, and R.C. Anderson. Onboard autonomous rock shape analysis for mars rovers. In *IEEE Aerospace Conference Proceedings*, March 2002.
- [23] M.S. Gilmore, R. Castaño, T. Mann, R.C. Anderson, E.D. Mjolsness, R. Manduchi, and R. Saunders. Strategies for autonomous rovers at mars. *Journal of Geophysical Research*, 105:223–229, 2000.

- [24] C.C. Gottlieb and H.E. Kreyszig. Texture descriptors based on co-occurrence matrices. *Computer Vision, Graphics and Image Processing*, 51:76–80, 1990.
- [25] J.P. Hyslip and L.E. Vallejo. Fractal analysis of the roughness and size distribution of granular materials. *Engineering Geology*, 48:231–244, 1997.
- [26] Jr. J.E. Dobkins and R.L. Folk. Shape development on Tahiti-nui. *Journal of Sedimentary Petrology*, 40:1167–1203, 1970.
- [27] B. Julesz. Visual pattern discrimination. *IRE Transactions on Information Theory*, 8:84–92, 1962.
- [28] B. Julesz. Textons, the elements of texture perception, and their interactions. *Nature*, 290:91–97, 1981.
- [29] J.M. Keller, S. Chen, and R.M. Crownover. Texture description and segmentation through fractal geometry. *Computer Vision, Graphics and Image Processing*, 45:150–166, 1989.
- [30] W.C. Krumbein. The effect of abrasion on the size, shape and roundness of rock fragments. *Journal of Geology*, 49:482–520, 1941.
- [31] W.C. Krumbein. Measurement and geological significance of shape and roundness of sedimentary particles. *Journal of Sedimentary Petrology*, 11(2):64–72, 1941.
- [32] V.F. Leavers. Use of the two-dimensional radon transform to generate a taxonomy of shape for the characterization of abrasive powder particles. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(12):1411–1423, December 2000.
- [33] L. Lepistö, I. Kunttu, J. Autio, , and A. Visa. Classification using non-homogenous textures and spectral imaging. In *WSCG Short Papers proceedings*, pages 82–86, February 2003.
- [34] L. Lepistö, I. Kunttu, J. Autio, and Visa A. Comparison of some content-based image retrieval systems with rock texture images. In *Proceedings of 10th Finnish AI Conference*, pages 156–163, December 2002.
- [35] L. Lepistö, I. Kunttu, J. Autio, and A. Visa. Classification method for colored natural textures using gabor filtering. In *Proceedings of 12th International Conference on Image Analysis and Processing*, pages 397–401, September 2003.
- [36] L. Lepistö, I. Kunttu, and A. Visa. Rock image classification using color features in gabor space. *Journal of Electronic Imaging*, 14(4), 2005.
- [37] T. Leung and J. Malik. Representing and recognizing the visual appearance of materials using three-dimensional textons. *International Journal of Computer Vision*, 43:29–44, 2001.
- [38] D. Martin, C. Fowlkes, and J. Malik. Learning to detect natural image boundaries using local brightness, color, and texture cues. *IEEE Transactions on Pattern Analysis Machine Intelligence*, 26(5):530549., 2004.
- [39] V. Mikli, H. Kaerdi, P. Kulu, and M. Besterci. Characterization of powder particle morphology. In *Proceedings of Estonian Academy of Sciences, Engineering*, volume 7, pages 22–34, 2001.
- [40] G. Mori. Guiding model search using segmentation. In *IEEE International Conference on Computer Vision*, 2005.
- [41] G. Mori, X. Ren, A. Efros, and J. Malik. Recovering human body configurations: Combining segmentation and recognition. In *IEEE Computer Vision and Pattern Recognition*, 2004.
- [42] R.J. Oakey, M. Green, P.A. Carling, M.W.E. Lee, D.A. Sear, and J. Warburton. Grain-shape analysis - a new method for determining representative particle shapes for populations of natural grains. *Journal of Sedimentary Research*, 75:1065–1073, 2005.
- [43] M. Partio, B. Cramariuc, M. Gabbouj, and A. Visa. Rock texture retrieval using gray level co-occurrence matrix. In *5th Nordic Signal Processing Symposium*, October 2002.
- [44] F. J. Pettijohn. *Sedimentary Rocks*. Harper, 2nd edition, 1957.
- [45] M. Peura and J. Iivarinen. Efficiency of simple shape descriptors. In *Proceedings of the Third International Workshop on Visual Form*, pages 443–451, May 1997.
- [46] M.C. Powers. A new roundness scale for sedimentary particles. *Journal of Sedimentary Petrology*, 23:117–119, 1953.

- [47] N.A. Riley. Projection sphericity. *Journal of Sedimentary Petrology*, 11(2):94–97, 1941.
- [48] J. Shi and J. Malik. Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8):888–905, 2000.
- [49] Z. B. Sozer. *Two-Dimensional Characterization of Topographies of Geomaterial Particles and Surfaces*. PhD thesis, Georgia Institute of Technology, March 2005.
- [50] D.R. Thompson, S. Niekum, T. Smith, and D. Wettergreen. Automatic detection and classification of geological features of interest. In *IEEE Aerospace Conference Proceedings*, March 2005.
- [51] D.R. Thompson, T. Smith, and D. Wettergreen. Data mining during rover traverse: From images to geologic signatures. In *8th International Symposium on Artificial Intelligence, Robotics and Automation in Space*, September 2005.
- [52] E. Trucco and A. Verri. *Introductory techniques for 3-D computer vision*. Prentice Hall, 1998.
- [53] M. Tuceryan and A.K. Jain. *Texture Analysis*, pages 207–248. World Scientific Publishing Co., 2nd edition, 1998.
- [54] M. Varma and A. Zisserman. A statistical approach to texture classification from single images. *International Journal of Computer Vision: Special Issue on Texture Analysis and Synthesis*, 62(1-2):61–81, 2005.
- [55] M. Varma and A. Zisserman. A statistical approach to material classification using image patch exemplars. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2007.
- [56] H. Wadell. Volume, shape, and roundness of rock particles. *Journal of Geology*, 40:443–451, 1932.
- [57] H. Wadell. Sphericity and roundness of rock particles. *Journal of Geology*, 41:306–309, 1933.
- [58] J. Wagner, G. Thomas, J. Glasgow, N. Cabrol, E. Grin, and R.C. Anderson. The accuracy of sediment size, shape, and distribution measurements from robotic geological images. *IEEE Transactions on Systems, Man and Cybernetics, Part A*, submitted January 2004.
- [59] L. Wang. Automatic identification of rocks in thin sections using texture analysis. *Mathematical Geology*, 27(7):847 – 865, October 1995.
- [60] Y. Weiss. Deriving intrinsic images from image sequences. In *Proceedings Eight IEEE International Conference on Computer Vision*, volume 2, pages 68–75, 2001.
- [61] H. Williams, F.J. Turner, and C.M. Gilbert. *Petrography: An Introduction to the Study of Rocks in Thin Sections*. W. H. Freeman and Company, 2nd edition, 1982.