Automatic Annotation of Planetary Surfaces with Geomorphic Labels

Soumya Ghosh and Tomasz F. Stepinski and Ricardo Vilalta

***Abstract*—We present a methodology for automatic geomorphic mapping of planetary surfaces that incorporates machine learning techniques. Our application transforms remotely sensed topographic data gathered by orbiting satellites into semantically meaningful maps of landforms; such maps are valuable research tools for planetary science. As topographic data becomes increasingly available, the ability to derive geomorphic maps efficiently is becoming essential. In our proposed framework, the mapping is achieved by means of scene segmentation followed by supervised classification of segments. The two mapping steps use different sets of features derived from digital elevation models (DEMs) of planetary surfaces; selection of appropriate features is discussed. Using a particular set of terrain attributes relevant to annotating cratered terrain on Mars, we investigate the design choices for both segmentation and classification components. The segmentation assessment includes K-means based agglomerative segmentation and watershed-based segmentation. The classification assessment includes three supervised learning algorithms: Naive Bayes, Bagging with decision trees, and Support Vector Machines (SVM); segments are classified into the following landforms: crater floors, crater walls (concave and convex), ridges (concave and convex) and inter-crater plains. The method is applied to six test sites on Mars. The analysis of the results shows that a combination of K-means based agglomerative segmentation and either SVM with a quadratic kernel or Bagging with C4.5 yields best maps. The presented framework can be applied without modification to generate geomorphic maps of sites on Earth.**

***Index Terms*—Digital topography, object recognition, segmentation, classification, Mars.**

I. Introduction

Remote sensing instruments onboard spacecrafts are providing increasingly large volumes of data on various aspects of Earth as well as surfaces of other planets. This data-rich environment challenges the scientific community to process, analyze, and ultimately turn into knowledge a significant portion of all collected data. In the context of planetary science, most data needs to be presented in the context of geomorphic maps [48] – basic staples of any planetary geology research. A geomorphic map is a categorical (thematic) map of topographical expressions (landforms). Traditionally, such maps are drawn by analysts through visual interpretations of images [42], [48].

In this paper we present a methodology for automating the process of geomorphic mapping. We are motivated by the slowness and expense of the traditional mapping process. If left to manual mapping, the percentage of planetary surface mapped to the level of detail permitted by an increased resolution of newly arriving data will continue to drop precipitously. Automating the mapping process can prevent this decline. We use topographic data as an input for machine mapping, because landforms can be better characterized by topographic features than by image features. Our method relies on segmentation-based classification of a landscape scene into constituent landforms. Such an approach yields results that, in their appearance and content, mimic manually derived maps.

S. Ghosh is with the University of Colorado, Boulder, CO 80309, USA (soumya.ghosh@colorado.edu)

T.F. Stepinski is with the Lunar and Planetary Institute, Houston, TX 77058-1113 USA (tom@lpi.usra.edu)

R. Vilalta is with the University of Houston, Houston, TX 77204-3010 USA (vilalta@cs.uh.edu)

A. Background

The overwhelming majority of previous work on automatic annotation of remotely sensed data was devoted to classification of multi-spectral images into different types of land covers. A standard approach to classification of raster data (images or digital elevation models or DEMs) is to use pixels as objects for classification [22], [23]. However, a pixel is not an ideal surface unit from the point of view of classification. Real life objects are characterized by a similarity of color and texture rather than a similarity of color alone. This observation led to the technique of segmentation-based classification wherein an image is first subdivided into meaningful objects that are subsequently classified [27], [45]. As a general rule, the segmentation-based approach produces better maps from images than the pixel-based approach [20]. The segmentation-based approach is even more attractive when working with topographic instead of imagery data. Pixels are just too small to constitute reasonable landscape elements; larger patches of surface better reflect local topographic expressions. This is important for supervised classification, where a domain expert must label each unit in a training set. Furthermore, geometric and statistical properties of landscape elements, such as shape, neighborhood properties, and statistics corresponding to terrain attributes calculated over an ensemble of member pixels, provide additional information that can be incorporated into classification.

The topic of auto-mapping landforms from topography has received little attention in the past, but with the introduction of segmentation-based classification methods, there is now a renewed interest in such mapping [13], [15]. However, published approaches rely on hand-made rules for classification and fail to take advantage of machine learning techniques.

B. Related Work

In the past, auto-mapping of planetary surfaces was restricted to identifying particular targets of interest, namely

impact craters [2], [3], [9], [10], [17], [21], [24], [25], [30], [34], [35], [47]. Machine crater identification could be considered an example of machine mapping, albeit a simple one, containing a single landform. In contrast, our work has been directed toward the automatic generation of geomorphic maps that contain several landforms. Our efforts have concentrated on planet Mars because Mars is the only planet besides Earth for which global topographic data is currently available [38]. In [8], [39] we explored pixel-based unsupervised learning for mapping landforms. Such an approach required a significant manual post-clustering processing phase to interpret the output clusters and produce a map amenable to interpretation by a planetary analyst. The analyst had no control over the classes of landforms expressed in the resulting map, and the interpreted classes lacked customary geomorphic meaning, thus limiting the practical use of the map. This occurs because a reasonable cluster derived under a proximity measure may not constitute a customary landform as perceived by a human analyst. Recognizing that appearance and content of automatically generated maps must mimic manually derived maps, we have started to explore the concatenation of a segmentation-based technique with supervised learning [40], [41] as a framework for geomorphic mapping.

C. Preface

In this paper we propose a unifying framework for annotating planetary surfaces with geomorphic labels. This framework is based on segmentation-based supervised learning. We utilize such framework for annotation of a particular surface - a cratered terrain on planet Mars. We assess and discuss the use of two different segmentation algorithms designed to procure small segments required for accurate classification. We also assess three supervised learning algorithms for the purpose of classifying segments into desired landform classes. Landform classification is challenging because different landforms of interest may be characterized by similar local terrain features but have different semantic meanings due to their spatial context. We address this issue by incorporating neighborhood information as segment-based features.

The rest of the paper is organized as follows. Section II describes the input data and feature selection step. In Section III we present our approach to scene segmentation; Section IV describes the classification techniques. Experimental results are shown and discussed in Section V, followed by conclusions and future research directions in Section VI.

II. Dataset and Pixel-Based Features

Our method is contingent on the availability of topographic data. In the near future, such data will become available for the Moon and planet Mercury, but presently Mars is the only planet (other than the Earth) for which global topography is available. Raw data used in this paper was gathered by the Mars Orbiter Laser Altimeter (MOLA) instrument [38] onboard the Mars Global Surveyor spacecraft. The data is organized in the form of a digital elevation model (DEM), a raster dataset where each pixel is assigned an elevation value.

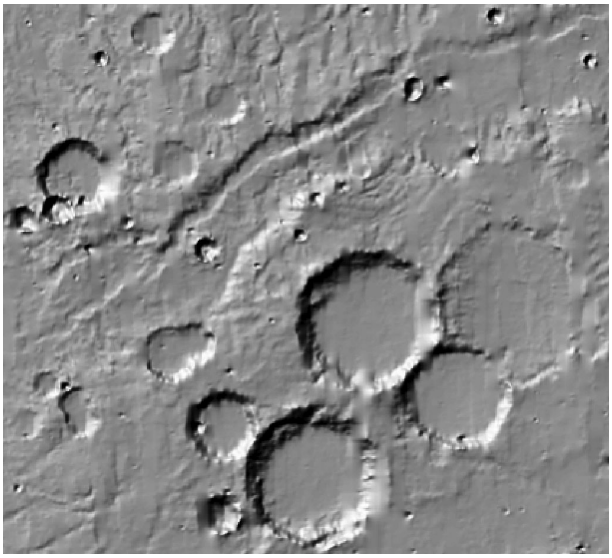

Fig. 1. Topography of the Tisia test site on Mars.

The global Martian DEM has a resolution of 128 pixels per degree (~ 460 meters at the equator).

We demonstrate the performance of our method by generating geomorphic maps for several sites on Mars referred to as Tisia, Al-Qahira, Dawes, Evros, Margaritifer, and Vichada. The Tisia site, centered at 46.13° E, 11.83° S, is our primary study area; it will be used to demonstrate the components of our framework and will serve as the training site. It was chosen as the primary site because it contains all the landforms of interest in a relatively small area. Moreover, the same site was previously studied in [39], [41], [44], so our present mapping can be directly compared with previous results. The DEM of Tisia site has dimensions of 385 rows and 424 columns and is shown in Fig. 1.

We generate a geomorphic map expressing the following landforms: crater floors, crater walls, ridges and inter-crater plains. Crater walls and ridges are further categorized into convex and concave depending on their curvature. Lower regions of walls and ridges have positive curvature and are termed convex. On the other hand, higher regions of walls and ridges have negative curvature and are termed concave. Overall, we aim at mapping six landform classes: crater floors, convex crater walls, concave crater walls, convex ridges, concave ridges, and inter-crater plateau. The choice of these particular landform classes stems from our interest in the automatic characterization of cratered surfaces. Other landforms could be chosen depending on the application of interest.

The DEM carries information about terrain elevation, $z(x, y)$ at the location (x, y) of each pixel. However, elevation is not a variable well suited to discriminate between the aforementioned landform classes. Instead we use the DEM to calculate other terrain attributes:

1) **Slope** — Slope, $s(x, y)$, defines the rate of maximum change of z on the 3×3 neighborhood of a pixel. A slope is a useful feature for delineating crater walls and ridges.
2) **Curvature** — The tangential curvature [28], $\kappa(x, y)$,

measures the change of slope angle in the direction of tangent to contour and is calculated analytically from a fourth-order polynomial fitted to a patch of surface composed of a 3×3 neighborhood centered on a focus pixel. Curvature is used for distinguishing between convex and concave crater walls and ridges. $\kappa > 0$ corresponds to convex topography, whereas $\kappa < 0$ corresponds to concave topography.

3) **Flood** — Flood, $f(x, y)$, is a binary variable such that pixels located inside topographic basins have $f(x, y) = 1$, and all other pixels have $f(x, y) = 0$. The flood variable is particularly useful to distinguish between two flat landforms, inter crater plateaus and crater floors, because crater floors are usually enclosed basins and plateaus are not.

A pixel-based feature vector, $v(x, y) = \{s, \kappa, f\}(x, y)$, describes the topography of the landscape at the level of an individual pixel. The raster of vectors v offers a mathematical representation of a site's landscape. Pixel-based feature vectors are used in the segmentation stage of our proposed method.

III. Segmentation

Segmentation is an intermediate data-preparation stage in our framework. The segmentation procedure subdivides the landscape into mutually exclusive and exhaustive segments containing pixels having approximately uniform feature vectors. These segments constitute topographic objects, which are classified and merged into coherent physical landforms by the supervised learning algorithms applied in the later stages of our framework. Raster segmentation has been the subject of intense study in the domain of image analysis. In the context of image analysis, a variety of techniques have been proposed [1], [4], [11], [14], [29], [37], [46], all of which could, in principle, be extended to landscape segmentation. However, requirements for an effective segmentation for the purpose of classification are different from those encountered in the field of computer vision. In particular, for the purpose of classification, it is desirable to have relatively small, approximately equal-sized segments (a requirement at odds with typical problems in image analysis). Having small segments eliminates the danger of a particularly large segment being misclassified, which would avoid producing a grossly incorrect map. Additionally, a misclassification of a small segment carries a small impact in the overall map accuracy. Moreover, having approximately equal-sized segments assures that statistics of terrain attributes are calculated from comparable ensembles of member pixels.

We investigate two different segmentation algorithms. The dividing algorithm splits the landscape on the basis of abrupt discontinuities in pixel-based feature vectors. The agglomerative algorithm initially treats each pixel as an individual segment. These initial segments are combined into larger segments as long as a user-defined criterion for the uniformity of constituent pixel-based feature vectors holds. Both algorithms use the same pixel-based feature vectors $v(x, y)$.

A. Watershed-Based Segmentation

Our first approach to segmenting a landscape incorporates the watershed transform [5] applied to a gray-scale image that encapsulates gradients of pixel-based feature vectors. This image is calculated using a (computationally simple) homogeneity measure H [18]. The homogeneity measure H is calculated using a moving square window of width $2L$ + 1 (where L is user-defined) applied over the raster of feature vectors $v(x, y)$. Consider a focal pixel (x_c, y_c) having a single feature (a component of a feature vector), for example, slope $s(x_c, y_c)$. For every pixel in the window we calculate a separation vector $\mathbf{d}_i = (x_i - x_c, y_i - y_c)$. From the separation vector we construct a gradient vector,

$$\mathbf{g}_i = (s(x_i, y_i) - s(x_c, y_c)) \times \frac{\mathbf{d}_i}{\|\mathbf{d}_i\|} \tag{1}$$

and we use gradient vectors calculated for all pixels in the window to calculate the homogeneity measure H,

$$H = \| \sum_{i=1}^{(2L+1)^2} \mathbf{g}_i \| \tag{2}$$

A pixel located in a region that is homogeneous with respect to s has a small value of H. On the other hand, a pixel located near a boundary between two regions characterized by different values of s has a large value of H. A raster constructed by calculating the values of H for all pixels in the landscape can be interpreted as a gray-scale image and is referred to as the H-image. We denote the H-image by H. White areas on H represent boundaries of homogeneous regions, whereas the dark areas represent the actual regions. The extension of the H-image concept to encapsulation of gradients in a field of multi-dimensional feature vectors is straightforward. In order to calculate H for our three dimensional feature vectors $v(x, y)$, we calculate the three individual H values separately for each pixel and combine them to obtain the overall value of H at that pixel:

$$H = \sqrt{w_s H_s^2 + w_\kappa H_\kappa^2 + w_f H_f^2} \tag{3}$$

where w_s, w_κ, and w_f are weights introduced to offset different numerical ranges of the attributes. An example of the H-image is given in Fig. 2A. The watershed transform of H results in an over segmentation of the H-image (and thus the landscape), an undesirable feature in computer vision, but a desirable feature in the context of segmentation-based classification.

B. K-Means Based Segmentation

Our second approach to segmenting a landscape is agglomerative and incorporates a contiguity-enhanced variant of the standard K-means clustering algorithm [43], which uses – in addition to terrain attributes – spatial coordinates of pixels as features. In this approach the pixel-based feature vector is given as $v = (s, \kappa, f, x, y)$. The additional spatial features (x and y) control the size of the segments while providing the resultant segments with very desirable geometric properties. For example, in areas where terrain features are approximately uniform, the local gradient of v is dominated by changes in x and y leading to the formation of round-shaped segments. On the other hand, in areas where change

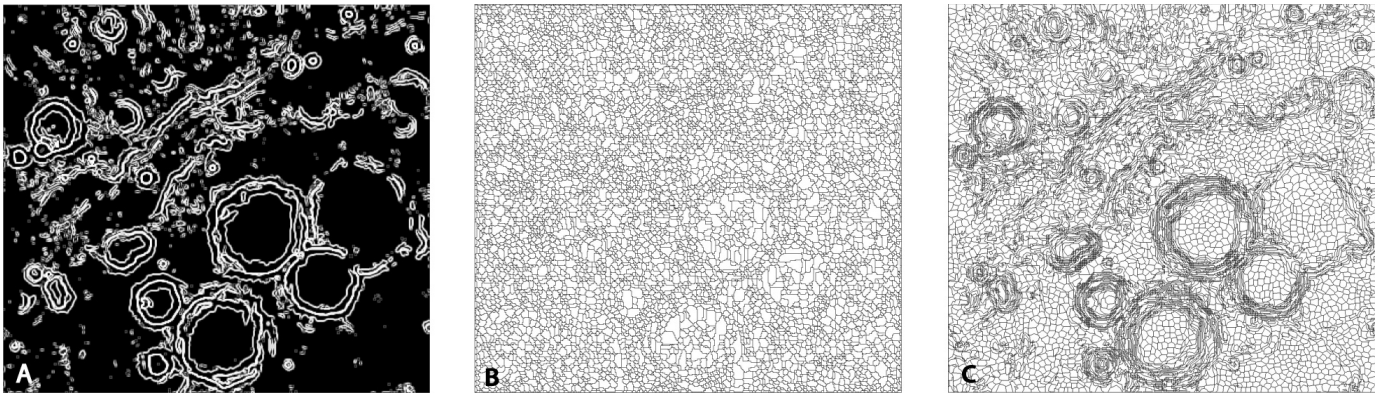

Fig. 2. Tisia site segmentation. (A) The H-image used for watershed based segmentation. (B) The watershed based segmentation. (C) The K-means based segmentation.

of terrain features dominates the local gradient of v, segments tend to exhibit an elongated shape in direction perpendicular to the gradient of the terrain-only sub-vector. These properties constitute additional knowledge that could be exploited by the classification module.

The actual segmentation invokes a simple K-means algorithm [33] applied to spatially-enriched, pixel-based feature vectors. The size of the segments is controlled by the value of k (which needs to be large to achieve over-segmentation). The resulting k clusters do not correspond to k single-connected spatial segments; instead each cluster may contain a number of segments. To derive the final segmentation (with $K > k$ segments) we assign a unique segment identifier to each subset of a cluster corresponding to a single spatially connected region.

C. Comparing Segmentation Results

We compare the two approaches to scene segmentation on the Tisia site. The watershed based algorithm with $L = 1$ produced 7708 segments with sizes ranging from 1 to 267 pixels. The mean segment size is 26 pixels and the standard variation of segment size is 22 pixels. In the k-means based algorithm we set $k = 5000$ resulting in 6593 single-connected segments having sizes ranging from 4 to 117 pixels. The mean segment size is 25 pixels and the standard variation of segment size is 16 pixels. Figs. 2B-C show the resultant segmentations. Notice the different character of the two kinds of segmentation with only the K-means based segmentation reflecting closely the site's topography (compare to Fig. 1).

In order to quantitatively compare the two segmentations we use compactness and isolation [43], two internal measures of clusters quality. Before we explain these measures we describe our notation. K stands for the total number of segments, with j being the segment index ($j = 1, \ldots, K$); v_i is a pixel-based feature vector in a segment with i being the pixel index; I_j is the set of pixels in segment j, with n_j being the size of that set; $N = \sum_j n_j$ is the total number of pixels in the site, and $c_j = (\sum_{i \in I_j} v_i)/n_j$ is the mean feature vector in segment j; $c = (\sum_j n_j c_j)/N$ is the mean feature vector calculated over all pixels in the site.

In calculating compactness and isolation we use only the non-spatial part of pixel-based feature vectors, $v = \{s, \kappa, f\}(x, y)$. The compactness V_j of a single segment is defined as the within-segment variance of its feature vectors:

$$V_j = \frac{1}{n_j} \sum_{i \in I_j} (v_i - c_j)^2 \tag{4}$$

On each segment we expect all feature vectors to be nearly identical, so that small values of V_j are desirable. We calculate compactness of the entire segmentation as a weighted average (based on size) of each individual segment compactness:

$$V = \frac{1}{N} \sum_{j=1}^{K} n_j V_j \tag{5}$$

A normalized version of V is obtained by dividing over the variance of the pixel-based feature vectors on the entire site:

$$V_0 = (\sum_{i=1}^{N} (v_i - c)^2)/N \tag{6}$$

Our final measure of segmentation compactness is defined as follows:

$$V_\star = V/V_0 \tag{7}$$

A segmentation with a smaller value of $V_\star$ is better inasmuch as it has segments that, on average, are more uniform.

The isolation S_j of a single segment is defined as the ratio of the feature-based distance between this segment and its closest neighbor, and the segment's feature-based size:

$$S_j = \sqrt{\frac{\min_l (c_l - c_j)^2}{(\sum_{i \in I_j} (v_j - c_j)^2)/n_j}} \tag{8}$$

where l indexes the neighbors of segment j. A large value indicates that the difference between a segment and its closest neighbor is significantly larger than the differences between individual pixels within the segment. Such segment is well isolated from its neighbors. We calculate the isolation of the

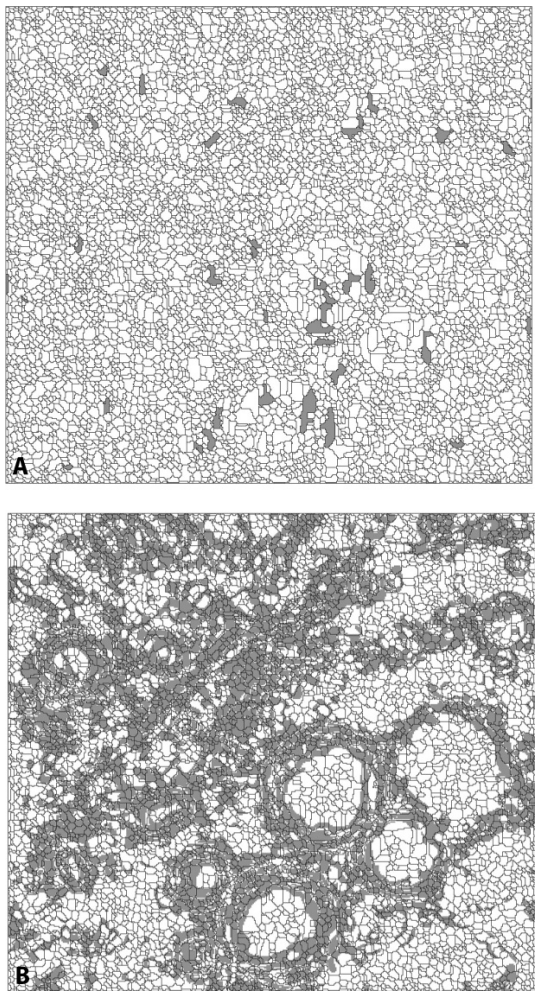

Fig. 3. Comparison of segments elongation in the Tisia site. Gray-colored segments indicate elongated segments with $SCI \geq 1.75$. (A) The watershed-based segmentation. (B) The K-means based segmentation.

entire segmentation as a weighted average (based on size) of each individual segments' isolation:

$$S = \frac{1}{N} \sum_{j=1}^{K} n_j S_j \tag{9}$$

A segmentation with a larger value of S is better inasmuch as it has segments that, on average, are more isolated from their neighbors. For the watershed-based segmentation of the Tisia site (Fig. 2B), $V_\star = 0.068$ and $S = 0.061$, whereas for the K-means based segmentation of the same site (Fig. 2C), $V_\star = 0.022$ and $S = 0.087$. These results indicate that K-means based segmentation is of significantly higher quality than the watershed based segmentation. Also, the K-means based segmentation has a preferable (flatter) distribution of segment sizes. Overall, the K-means based segmentation shows better performance with the only disadvantage of having higher computational cost [1].

[1]In the computer vision community K-means is generally perceived to be faster and more efficient than state of the art segmentation routines. However, this is only true for small values of k. In our application where $k \sim 10^3 - 10^4$, k-means turns out to be computationally expensive.

IV. Classification

The classification module assigns a label (landform designation) to unlabeled segments based on patterns learned from examples. In our framework, the subjects of classification are segment-based feature vectors associated with individual segments. It is important to stress that these segment-based feature vectors are not the same as the pixel-based feature vectors used for segmentation. First, the location information, contained in spatial coordinates for the K-means based segmentation, is no longer necessary or desirable in the segment-based vectors; the landform designation should not depend on a specific location of a landscape element. The major components of the segment-based features are terrain parameters averaged over an ensemble of member pixels (i.e. vectors c_j). Additional segment-based features are calculated from the shape of the segments and properties of neighboring segments. The spatial features describe the segment itself (i.e., its geometrical and neighboring properties).

A. Segment-Based Features

Each segment is represented by a combination of physical and spatial features. Physical features comprise segmentation features averaged over the constituent pixels of the segments; they are derived from components of the vector c_j: $\bar{s_j}$, $\bar{\kappa_j}$, and $\bar{f_j}$, where the bar indicates averages over the segment's constituent pixels.

Spatial features are obtained using the segment's shape measure and the neighborhood context measure. The shape measure is computed in terms of the **S**hape **C**omplexity **I**ndex (SCI) [16] for each segment as follows:

$$SCI = \frac{P}{2\pi r}; \quad r = \sqrt{\frac{A}{\pi}} \tag{10}$$

Here P is the perimeter of the segment's boundary, A is the area of the segment, and r is the radius of a circle with the same surface area as the segment. SCI is a measure of segment circularity. The closer the value of SCI to 1.0, the more circular the object; on the other hand, thin ring-like shapes tend to have SCI values of 2.5 and higher. The perimeter P and area A of the segments are calculated using Minkowski functionals [26].

We expect that the SCI to be a valuable feature for landform identification. Note that this feature is only available in the segmentation-based classification approach, and not available in the standard pixel-based classification approach. In order to investigate the importance of the SCI in delineating landforms, we have produce a binary image $B_{\rm SCI}$ for the Tisia site, according to the following transformation:

$$B_{\rm SCI}(j) = \begin{cases} B_{\rm SCI}(j) = 1 & \text{for } SCI(j) \geq SCI_{\rm th} \\ B_{\rm SCI}(j) = 0 & \text{for } SCI(j) < SCI_{\rm th} \end{cases} \tag{11}$$

where, $j = 1, ..., K$ indexes the segments, and $SCI_{\rm th}$ is a threshold value that decides when segments are considered elongated. Figs. 3A-B show the images $B_{\rm SCI}$ for the watershed based and the k-means based segmentations of the Tisia site using $SCI_{\rm th} = 1.75$. There is a striking difference between the two images. Whereas in the K-means based segmentation the segments corresponding to ridges and crater walls are

TABLE I
NUMBER OF LABELED SEGMENTS PER CLASS

Segmentation	Classes						Labeled	Total
	1	2	3	4	5	6		
watershed	788	291	156	156	57	50	1498	7708
K-means	345	243	126	67	18	30	829	6593

elongated and have $SCI > SCI_{\rm th}$, no such pattern is observed in the watershed based segmentation, indicating that most watershed-derived segments are fairly circular. This difference stems from the way the two algorithms define segments. In the watershed algorithm, segments are defined by the smallest "depressions" in the H-image, essentially by the noise in the H-image, which leads to relatively circular shapes. On the other hand, the contiguity-enhanced K-means algorithm is especially designed (see Section 3.2) to produce segments with shapes that reflect gradients of physical features. Thus, SCI is helpful in distinguishing landforms only when used in conjunction with the K-means based segmentation.

As mentioned in Section 1.3, the automatic classification of landforms is challenging because different landforms may be characterized by very similar terrain attributes, but different spatial context. For instance, segments making up craters' walls and segments constituting ridges not associated with craters may have similar values of slope, curvature, and flood, but located in different spatial contexts. Our method takes into consideration spatial context by means of neighborhood context measures. Ideally, we would like to know classes of segment's neighbors to establish its spatial context, but such information is not available prior to classification. However, even for unlabeled segments, a preliminary categorization of segments into low, medium, and high slope categories is possible on the basis of statistics of the values of $\bar{s_j}$, $j = 1, \ldots, K$. Such categorization is used to calculate the neighborhood property of a segment j, $\{a_1^s,\ a_2^s,\ a_3^s\}_j$, where $a_l^s,\ l = 1, 2, 3$ is the percentage of the segment boundary adjacent to neighbors belonging to slope category l. Similar neighborhood properties, $\{a_1^\kappa,\ a_2^\kappa,\ a_3^\kappa\}_j$, $\{a_1^f,\ a_2^f,\ a_3^f\}_j$ are calculated on the basis of curvature and flood values, yielding a total of nine features corresponding to the spatial context of a given segment. Segment-based feature vectors, denoted by u, are defined as follows when k-means based segmentation is used:

$$\begin{aligned} u \ = \ & \{\bar{s},\ \bar{\kappa},\ \bar{f},\ SCI, a_1^s,\ a_2^s,\ a_3^s,\ a_1^\kappa, \\ & a_2^\kappa,\ a_3^\kappa,\ a_1^f,\ a_2^f,\ a_3^f\} \end{aligned} \quad (12)$$

The definition is slightly modified when watershed based segmentation is used by removing the SCI component.

B. Labeled Training Set

The labeled (training) set of segments was generated by manually labeling 30% (by surface area) of the Tisia site into six classes corresponding to six landforms as described in Section 2: inter-crater plateau (class 1), crater floors (class 2), convex crater walls (class 3), concave crater walls (class 4), convex ridges (class 5), and concave ridges (class 6). Fig. 4A shows the labeled part of the Tisia site. Depending on the segmentation used, the labeled area results in a different number of labeled segments. Table I provides details on the number of labeled segments.

TABLE II
ACCURACY FOR THE THREE CLASSIFIERS AND TWO SEGMENTATIONS TECHNIQUES (STANDARD DEVIATIONS ARE GIVEN IN PARENTHESES). AN ASTERISK ON AN ENTRY REPRESENTS STATISTICALLY SIGNIFICANT IMPROVEMENT OVER THE NAIVE BAYES CLASSIFIER (AT 95% CONFIDENCE LIMITS USING A T-STUDENT DISTRIBUTION).

Dataset	Naive Bayes	Bagging with C4.5	SVM (quadratic kernel)
K-means	88.65(3.17)	$90.95(2.57)^*$	$91.06(2.50)^*$
watershed	85.27(2.51)	$90.28(1.95)^*$	$91.30(1.58)^*$

C. Classifiers

We applied different learning algorithms in our test sites for segment classification and to generate geomorphic maps. First we tried a simple classifier as a baseline for comparison to determine if a limited family of models could perform satisfactorily in our particular domain. Naive Bayes belongs to the class of generative Bayesian classifiers that estimate the posterior probability of the class label given a feature vector representing the instance (segment) using Bayes' theorem [19]. The computation of the likelihood is simplified by assuming feature independence given the class label. The independence assumption is false in our case as all the features are derived from a single source of information – the site's DEM.

Our second classifier imposes a flexible model over the feature space and was used to test the importance of reducing bias by employing a rich family of models. Support Vector Machines (SVM) is a statistical learning algorithm that works by finding an optimal hyper-plane in a (transformed) feature space [6]. The optimal hyper-plane maximizes the separation between classes. SVM exploits local data patterns and has been found to be effective in spatial data mining applications [36].

We finally employed a classifier that reduces variance through model combination. The goal was to explore the importance of reducing generalization error incurred through model variability. Bagging is an ensemble learning algorithm [7], [12]; it generates multiple models by running a single learning algorithm multiple times over bootstrapped samples of the training set. The final class label is the result of voting over the contributing models (one from each bootstrap sample). Bagging is known to work well for complex datasets and is particularly attractive when the training set is noisy [12] as is the case in our application. We use a decision tree (C4.5) as the base learner.

We employed the algorithms above as implemented in the software package WEKA [49]. We used default parameters because our empirical assessment serves simply to illustrate the potential application of this framework. Nevertheless, we performed a parameter search over the Gaussian, linear, quadratic and cubic kernels in the case of Support Vector Machines at different values of the slack variable. The best results were obtained using a quadratic kernel with the slack value ranging between $C = 1$ and 2.

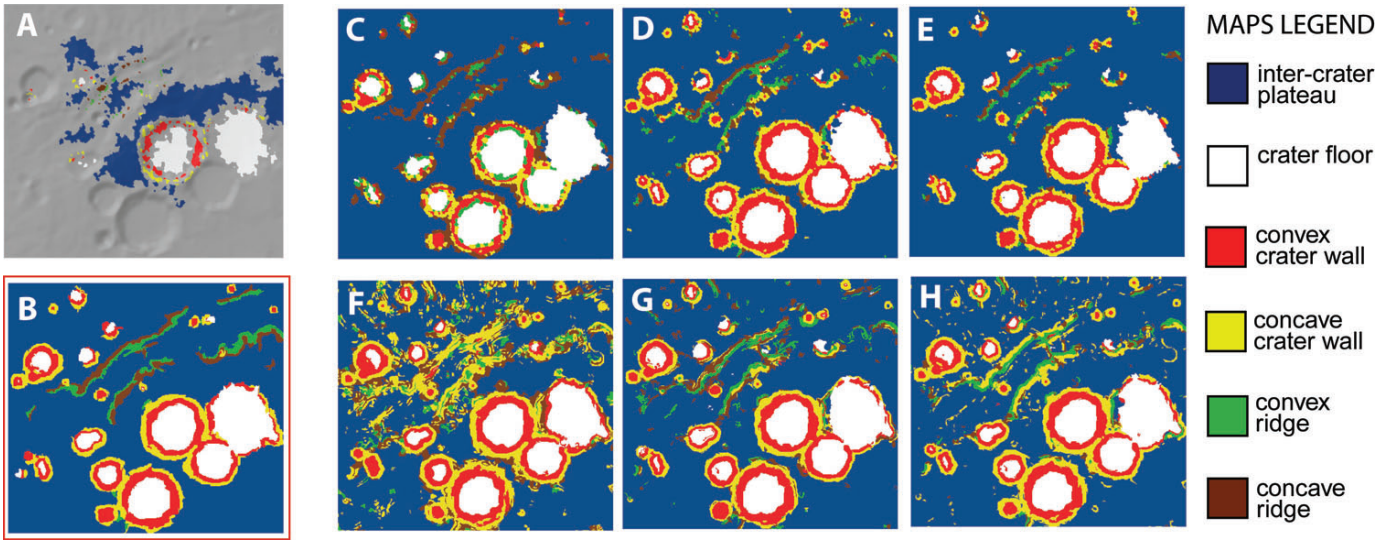

Fig. 4. Geomorphic maps of the Tisia site generated by different segmentation/classification methods. (A) Labeled regions; different landform classes are coded by different colors as given by the legend. Unlabeled regions are displayed in gray. (B) The analyst-drawn map ("ground truth"). Maps based on the watershed segmentation (C – Naive Bayes, D – Bagging, E – SVM). Maps based on the K-means segmentation (F – Naive Bayes, G – Bagging, H – SVM).

Accuracy of the three classifiers described above was computed using 10 runs of 10-fold cross-validation on the Tisia site training set (see Section 4.2). Accuracy is the percentage of pixels that are labeled in agreement with the analyst-drawn map. Results are shown in Table II. Both Bagging with C4.5 and SVM show significant improvement in accuracy compared with Naive Bayes, which points to the importance of imposing a rich family of models by the classifier in our domain. There is no significant difference in performance between Bagging and SVM.

V. Results

The segmentation-based classification technique described in Sections 3 and 4 was applied to six sites on Mars. For each site we have calculated the watershed-based and the K-means based segmentations. Each segmentation solution was classified using three learning algorithms (Naive Bayes, Bagging with C4.5 and SVM) trained using the same labeled training set (a subset of Tisia site) described in Section 4.2.

A. Tisia site

For our primary test site (Tisia) we have generated six maps shown in Figs. 4C-E and Figs. 4F-H. Fig. 4B shows the hand-labeled "ground truth" map of Tisia. This is how a typical domain expert would map the six landforms in this site. A human-produced map does not really constitute a ground truth because an analyst is likely to draw an idealized map that misses details and projects a human conceptualization of the entire landscape, even if it contradicts local measurements. Maps based on the watershed segmentation (Figs. 4C-E) have a "simple" look as they lack small-scale details, whereas maps based on the K-means segmentation (Figs. 4F-H) look "busy" as they exhibit more small-scale details. On the basis of a casual visual inspection one could conclude that maps based on watershed segmentation look more like the analyst-drawn map (Fig. 4B) than the maps based on K-means segmentation; this is despite the higher quality of K-means segmentation (Section 3.3), and the fact that segment-based feature vectors used to classify K-means generated segments carry relevant information. However, closer inspection of the generated maps shows that maps based on K-means segmentation correctly reflect some small-scale details that are absent from the watershed segmentation and the analyst-drawn map. Overall, maps generated by Naive Bayes are inaccurate and inferior to maps generated by Bagging and SVM. Maps generated by Bagging and SVM have different character; the SVM-generated map shows more detail, whereas the Bagging-generated map shows better discrimination between crater walls and ridges.

Detailed information about the six maps calculated for the Tisia test site are given in Table III. The two horizontal sections of this table report on results obtained using the watershed based segmentation and the K-means based segmentation, respectively. Each section has six subsections. The first subsection gives the names of the landforms. The second subsection gives the number of segments in each landform class. The remaining four subsections report on average values of different terrain parameters in all six landform classes. These values are normalized (a class with highest value receives a score of 100). For example, in a map produced by SVM on the basis of the watershed segmentation, the convex crater wall class is characterized by the highest values of slope (a score of 100). The slopes in the concave crater wall class are on average only 80% as steep (score 80), the convex and concave ridges classes are both on average 41% as steep, the plateau class is 3%, and the crater floor class is the least steep, on average only 1% as steep as the convex crater wall class. For the curvature attribute, positive values refer to convex regions, while negative values refer to concave regions.

Table IV gives accuracy rates for maps of the Tisia site. Disregarding maps produced by Naive Bayes, accuracy rates

TABLE III
PROPERTIES OF TOPOGRAPHIC LANDFORM CLASSES AS GIVEN BY DIFFERENT CLASSIFIERS APPLIED TO THE TISIA SITE. FOR EACH CLASSIFICATION SOLUTION (MAP) AND TERRAIN ATTRIBUTE A COLUMN GIVES AVERAGE MAGNITUDE OF THIS ATTRIBUTE IN A GIVEN CLASS (RANGING FROM ± 100 TO 0). **NB** – NAIVE BAYES, **B** – BAGGING WITH C4.5, **SVM** – SUPPORT VECTOR MACHINES.

watershed based segmentation															
class	**segments**			**area**			**slope**			**curvature**			**flood**		
	NB	**B**	**SVM**	**NB**	**B**	**SVM**	**NB**	**B**	**SVM**	**NB**	**B**	**SVM**	**NB**	**B**	**SVM**
plateau	5588	5354	5809	100	100	100	1	3	3	-1	-1	-1	1	2	3
floor	569	464	476	15	14	15	1	1	1	10	6	6	100	100	100
cvx. wall	384	687	573	3	10	9	100	100	100	79	66	68	92	92	94
con. wall	324	675	572	4	10	8	83	79	80	-84	-100	-100	38	7	11
cvx. ridge	210	258	136	1	1	1	46	35	41	100	46	54	79	2	6
con. ridge	616	270	142	7	1	2	36	40	41	-70	-35	-49	1	1	1
K**-means based segmentation**															
class	**segments**			**area**			**slope**			**curvature**			**flood**		
	NB	**B**	**SVM**	**NB**	**B**	**SVM**	**NB**	**B**	**SVM**	**NB**	**B**	**SVM**	**NB**	**B**	**SVM**
plateau	2660	3811	3613	100	100	100	3	1	3	-1	-1	1	1	1	2
floor	463	614	519	14	15	15	1	1	1	7	6	3	100	100	100
cvx. wall	793	739	754	13	9	11	100	100	100	100	52	82	97	91	95
con. wall	1846	739	1305	30	7	18	58	70	62	-70	-100	-100	8	4	8
cvx. ridge	238	272	254	1	1	1	20	25	36	10	43	55	7	8	7
con. ridge	593	355	148	11	2	1	23	29	31	-5	-16	-2	1	1	1

TABLE IV
PERFORMANCE MEASUREMENTS FOR MAPS OF THE TISIA SITE GENERATED BY DIFFERENT CLASSIFIERS. THE ENTRIES FOR INDIVIDUAL LANDFORMS ARE (PRECISION / RECALL). **NB** – NAIVE BAYES, **B** – BAGGING WITH C4.5, **SVM** – SUPPORT VECTOR MACHINES.

watershed based segmentation							
classifier	**overall accuracy**	**plateau**	**floor**	**cvx. wall**	**con. wall**	**cvx. ridge**	**con ridge**
NB	84.67	0.94 / 0.97	0.93 / 0.93	0.87 / 0.39	0.57 / 0.37	0.12 / 0.12	0.20 / 0.51
B	89.31	0.95 / 0.96	0.97 / 0.87	0.80 / 0.83	0.65 / 0.75	0.51 / 0.45	0.38 / 0.37
SVM	89.71	0.92 / 0.98	0.96 / 0.90	0.87 / 0.75	0.72 / 0.70	0.62 / 0.29	0.48 / 0.29
K**-means based segmentation**							
classifier	**overall accuracy**	**plateau**	**floor**	**cvx. wall**	**con. wall**	**cvx. ridge**	**con ridge**
NB	74.11	0.99 / 0.76	0.99 / 0.80	0.70 / 0.78	0.32 / 0.83	0.14 / 0.19	0.08 / 0.25
B	87.42	0.96 / 0.92	0.94 / 0.92	0.80 / 0.85	0.67 / 0.72	0.34 / 0.47	0.25 / 0.40
SVM	86.10	0.97 / 0.91	0.98 / 0.87	0.77 / 0.79	0.46 / 0.83	0.46 / 0.54	0.18 / 0.12

are above 86%. Note that maps based on the watershed segmentation have slightly higher rates than maps based on the K-means segmentation. This reflects the watershed segmentation simpler character, more in line with the analyst drawing. The table also shows precision and recall rates for six landform classes. Results show that inter-crater plateau, crater floor, and convex crater walls landforms are designated with high accuracy. Concave crater walls are detected with less accuracy, and ridges are essentially difficult to identify correctly. This is because local ridges look like crater walls, even though they are different landforms in the context of the entire landscape.

B. Other sites

The remaining five test sites, Vichada, Al-Qahira, Dawes, Evros, and Margaritifer, represent the same type of Martian surface as the Tisia site. The same six landform classes are present; one would then expect that models generated from the Tisia site would be able to map these sites accurately. In general, maps of these sites have the same relative character as maps from Tisia site. Fig. 5 shows a sample of the results. The top row in this figure shows the topography of each site. The topography helps to visualize the landscape and serves as a guide for a visual assessment of map quality. The middle row shows maps generated by the SVM classifier applied to the watershed segmentation, and the bottom row shows maps generated by the SVM classifier applied to the K-means segmentation. As in the case of the Tisia site, maps originating from the watershed segmentation lack small-scale details, whereas maps originating from the K-means segmentation show more details, but tend to generate more misclassifications.

The Vichada site was hand-labeled by an analyst. Table V shows accuracy performance for maps of this site. Maps corresponding to watershed segmentation have accuracy rates only slightly lower than those of Tisia site. However, these maps are much less useful than Tisia because they miss most ridges. The Vichada site is dominated by the plateau class, identified correctly during prediction (which explains the high accuracy rates). One gains more insight when evaluating precision/recall rates for individual landforms; such evaluation shows poor performance on all classes except the plateau class. Maps originating from the K-means segmentation have precision/recall rates similar to those recorded for the Tisia site.

VI. SUMMARY AND CONCLUSIONS

In this paper we have proposed a machine learning-based method for annotating planetary surfaces with geomorphic

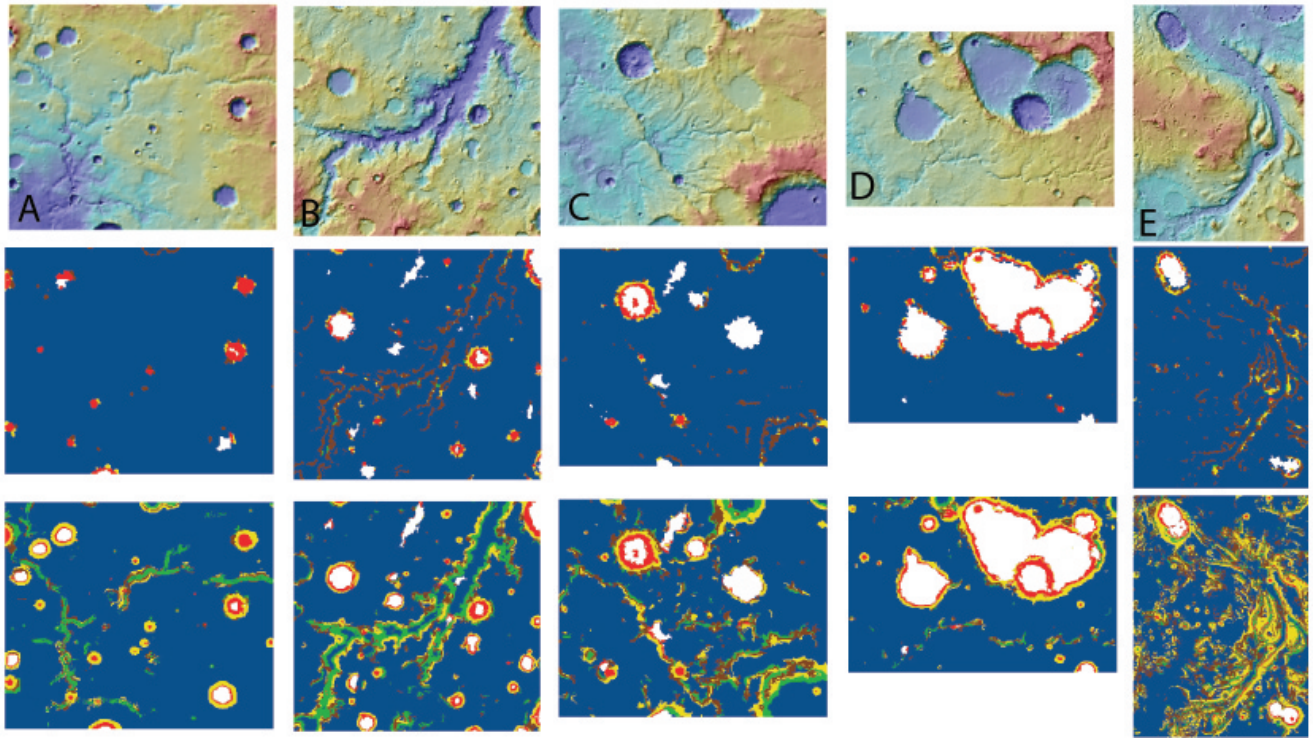

Fig. 5. Automatically generated six-landform geomorphic maps of Vichada (A), Al-Qahira (B), Dawes (C), Evros (D), and Margaritifer (E) sites using the SVM classifier. The top row shows sites' topography, the middle row shows maps originating from the watershed segmentation, and the bottom row shows the maps originating from the K-means segmentation. For the map legend see Fig. 4.

TABLE V
ACCURACY OF MAPPING INDIVIDUAL LANDFORMS IN THE VICHADA SITE USING DIFFERENT CLASSIFIERS. THE ENTRIES FOR LANDFORMS ARE (PRECISION / RECALL). **NB** – NAIVE BAYES, **B** – BAGGING WITH C4.5, **SVM** – SUPPORT VECTOR MACHINES.

watershed based segmentation							
classifier	**overall accuracy**	**plateau**	**floor**	**cvx. wall**	**con. wall**	**cvx. ridge**	**con ridge**
NB	83.05	0.86 / 0.99	0.40 / 0.30	0.73 / 0.21	0.41 / 0.06	0.001 / 0.0003	0.025 / 0.007
B	85.03	0.87 / 0.99	0.49 / 0.25	0.66 / 0.60	0.60 / 0.31	0.35 / 0.03	0.07 / 0.003
SVM	83.79	0.85 / 0.99	0.44 / 0.23	0.66 / 0.37	0.87 / 0.07	0.0 / 0.0	0.02 / 0.002
K-means based segmentation							
classifier	**overall accuracy**	**plateau**	**floor**	**cvx. wall**	**con. wall**	**cvx. ridge**	**con ridge**
NB	39.60	0.97 / 0.36	0.0 / 0.0	0.62 / 0.73	0.31 / 0.09	0.81 / 0.11	0.67 / 0.25
B	69.20	0.94 / 0.75	0.71 / 0.78	0.54 / 0.69	0.73 / 0.46	0.10 / 0.84	0.09 / 0.008
SVM	86.00	0.92 / 0.95	0.60 / 0.82	0.80 / 0.44	0.60 / 0.66	0.41 / 0.56	0.25 / 0.08

labels (i.e., for machine generation of geomorphic maps). This is a challenging problem for the following reasons. First, landscapes are more difficult to label than multi-spectral images because their characterization is based on fewer features. A multi-spectral image has at least six channels (colors), whereas landform characterization tends to be limited (our study derives three topographic features). Second, space-related data exhibits less quality than terrestrial data. A Martian DEM is much coarser and noisier than terrestrial DEMs. High quality DEMs facilitate extraction of additional features (such as several types of surface curvature). Moreover, terrestrial topographic data can often be supplemented by other datasets, such as, for example, hydrological data. Such additional datasets are not available in space applications. Developing an application for mapping planetary surfaces, therefore, presents challenges not found in the terrestrial domain. Our approach is based on a segmentation step preceding classification; this is because segments offer the opportunity to create additional features, taking into account spatial context (albeit only local context). This work is intended to serve as a guideline for future work related to automated planetary mapping. It is important to add that the particular landform classes presented in this paper are examples of geomorphic studies of cratered terrain on planet Mars. In order to apply our framework to generate maps of different landforms of interest on Mars, Earth or other planets, different classes must be identified by a domain expert (resulting in different training sets).

As a result of our assessment of different segmentation and classification algorithms, we offer the following design choices for future machine mapping applications. For a rough map

that needs to be generated rather quickly, a combination of watershed based segmentation with either Bagging or SVM is preferable. An application based on such a combination may be useful in cases where maps need to be generated onboard a spacecraft. For a more detailed map, a combination of K-means based segmentation with either Bagging or SVM is recommended. Since the K-means segmentation is slow, applications based on such a combination may be more appropriate for an off-line analysis. Bagging is more sensitive to small local changes in topographic data, and is preferable in applications where high sensitivity is desired; an example is found in maps designed to detect valley networks on Mars. SVM is less sensitive to small changes in topography and produces smoother maps (this is not so evident on the Tisia site, but quite evident on maps from other sites). SVM is preferable in applications where smoother maps are desirable.

Our study also shows the necessity of using a representative training set. Our six test sites are from roughly similar terrains (same geological unit); one would expect that our training set, extracted from a single site, should yield similar results on other sites. This was not the case. A comparison of Tables IV and V indicates that the Tisia site was mapped with higher accuracy than the Vichada site. Also, visual inspection of Fig. 5 shows that the Margaritifer site is poorly mapped, even though it bears topographic similarity to the Al-Qahira site (the former resulting in a good map). Although different sites may have similar landforms, particular characterizations may exhibit different ranges of terrain parameters. As an example, ridges in Al-Qahira and Margaritifer sites look similar, but exhibit different slope magnitudes. Hence, in future applications, we advocate flagging segments labeled with low confidence for additional manual labeling.

We envision two types of spin-off applications. First is an off-line automated data analysis module. Frequently, geological analysis consists of a descriptive comparison among maps (of two or more sites) in order to identify similarities and/or differences in the presence and character of selected landforms. Because our methodology produces digital thematic maps of landforms, such maps can be subjected to machine-assisted comparison (see, for example [31], [32]). Second is an onboard application that generates geomorphic maps for the purpose of data compression. Such maps can be cheaply transmitted to Earth to assist ground controllers in decision making. Future work will address faster algorithms for quality segmentation (to replace the slow K-means algorithm), development of additional topography-derived features, incorporating imagery data to add texture information, and a robust approach to incorporate context information (e.g., using Markov Random Fields).

Acknowledgments. This work was supported by the National Science Foundation under Grants IIS-0431130, IIS-448542, IIS-0430208, and by NASA under grant NNG06GE57G. A portion of this research was conducted at the Lunar and Planetary Institute, which is operated by the USRA under contract CAN-NCC5-679 with NASA. This is LPI Contribution No. 1433.

REFERENCES

[1] R. Adams and L. Bischof. Seeded region growing. *IEEE Transactions on Pattern Analysis Machine Intelligence*, 16(6):641–647, 1994.

[2] L. Bandeira, J. Saraiva, and P. Pina. Impact crater recognition on mars based on a probability volume created by template matching. *IEEE Transactions on Geoscience and Remote Sensing*, 45(12):4008–4015, December 2007.

[3] T. Barata, E. I. Alves, J. Saraiva, and P. Pina. Automatic recognition of impact craters on the surface of mars. In *Image Analysis and Recognition*, volume 3212 of *Lecture Notes in Computer Science*, pages 489–496. Springer Berlin / Heidelberg, 2004.

[4] S. Belongie, C. Carson, H. Greenspan, and J. Malik. Color- and texture-based image segmentation using EM and its application to content-based image retrieval. In *Proceedings of the Sixth International Conference on Computer Vision*, 1998.

[5] S. Beucher. The watershed transformation applied to image segmentation. *Scanning Microscopy International*, (6):299–314, 1992.

[6] B. E. Boser, I. Guyon, and V. Vapnik. A training algorithm for optimal margin classifiers. In *Computational Learning Theory*, pages 144–152, 1992.

[7] L. Breiman. Bagging predictors. *Machine Learning*, 24(2):123–140, 1996.

[8] B. D. Bue and T. F. Stepinski. Machine detection of martian impact craters from digital topography data. *IEEE Transactions on Geoscience and Remote Sensing*, 45(1):265–274, 2007.

[9] M. C. Burl, T. Stough, W. Colwell, E. B. Bierhaus, W. J. Merline, and C. Chapman. Automated detection of craters and other geological features. In *6th International Symposium on Artificial Intelligence, Robotics and Automation in Space Montreal, Canada*, June 2001.

[10] Y. Cheng, A. E. Johnson, L. H. Matthies, and C. F. Olson. Optical landmark detection for spacecraft navigation. In *AAS/AIAA Astrodynamics Specialist Conference Puerto Rico*, Feb 2003.

[11] Y. Deng and B. S. Manjunath. Unsupervised segmentation of color-texture regions in images and video. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(8):800–810, 2001.

[12] T. G. Dietterich. Ensemble methods in machine learning. *Lecture Notes in Computer Science*, 1857:1–15, 2000.

[13] L. Dragut and T. Blaschke. Automated classification of landform elements using object-based image analysis. *Geomorphology*, 81:330 – 344, 2006.

[14] H. Feng, D. A. Castanon, and W. C. Karl. A curve evolution approach for image segmentation using adaptive flows. *International Conference on Computer Vision*, 02:494, 2001.

[15] A. L. Gallant, D. D. Brown, and R. M. Hoffer. Automated mapping of hammond's landforms. *IEEE Geoscience and Remote Sensing Letters*, 2(4):384–288, 2005.

[16] T. Hengl. *Pedometric mapping: bridging the gaps between conventional and pedometric approaches*. PhD thesis, Wageningen University, 2003.

[17] R. Honda, Y. Iijima, and O. Konishi. Mining of topographic feature from heterogeneous imagery and its application to lunar craters. In *Progress in Discovery Science, Final Report of the Japanese Discovery Science Project*, page 395407, London, UK, 2002. Springer-Verlag.

[18] F. Jing, M. Li, H. Zhang, and B. Zhang. Unsupervised image segmentation using local homogeneity analysis. In *IEEE International Symposium on Circuits and Systems*, 2003.

[19] G. H. John and P. Langley. Estimating continuous distributions in bayesian classifiers. In *Eleventh Conference on Uncertainty in Artificial Intelligence, San Mateo*, pages 338–345, 1995.

[20] J. John H. Lowry, R. Ramsey, K. Boykin, D. Bradford, P. Comer, S. Falzarano, W. K. J. Kirby, L. Langs, J. Prior-Magee, G. Manis, L. O'Brien, K. Pohs, W. Rieth, T. Sajwaj, S. Schrader, K. A. Thomas, D. S. K. Schulz, B. Thompson, C. Wallace, C. Velasquez, E. Waller, and B. Wolk. Final report on landcover mapping methods. Technical report, The southwest regional gap analysis project, october 2005.

[21] J. Kim, J. Muller, S. van Gasselt, J. Morley, G. Neukum, and t. Team. Automated crater detection, a new tool for mars cartography and chronology. *Photogrammetric Engineering and Remote Sensing*, 71(10):1205–1217, October 2005.

[22] D. Landgrebe. The Evolution of Landsat Data Analysis. In *Photogrammetric Engineering and Remote Sensing*, pages 859–867, 1997.

[23] D. Landgrebe. *Information Extraction Principles and Methods for Multispectral and Hyperspectral Image Data*, chapter 1. Information Processing for Remote Sensing,World Scientific Publishing Co, 1999.

[24] B. Leroy, G. Medioni, E. Johnson, and L. Matthies. Crater detection for autonomous landing on asteroids. *Image and Vision Computing*, 19(11):787 – 792, September 2001.

[25] R. Martins, P. Pina, J. S. Marques, and M. Silveira. Crater Detection by a Boosting Approach. *IEEE Geoscience and Remote Sensing Letters*, 6:127–131, Jan. 2009.
[26] K. Michielsen and H. De Raedt. Integral-geometry morphological image analysis. *Physics Reports*, 347:461–538, jul 2001.
[27] G. Miliaresis and N. Kokkas. Segmentation and object-based classification for the extraction of the building class from lidar dems. *Comput. Geosci.*, 33(8):1076–1087, 2007.
[28] H. Mitasova and L. Mitas. Interpolation by regularized spline with tension: I. theory and implementation. *Theory and implementation, Mathematical Geology*, 25(6):641–655, 1993.
[29] R. Nock and F. Nielsen. Statistical region merging. *PAMI*, 26(11):1452–1458, November 2004.
[30] C. Plesko, S. Brumby, E. Asphaug, D. Chamberlain, and T. Engel. Automatic crater counts on mars. In S. Mackwell and E. Stansbery, editors, *Lunar and Planetary Institute Conference Abstracts*, pages 1935–1936, Mar. 2004.
[31] T. K. Remmel and F. Csillag. Spectra of coincidences between spatial data sets: Toward hierarchical inferential comparisons of categorical maps. In *GeoComputation Conference*, 2005.
[32] T. K. Remmel and F. Csillag. Mutual Information Spectra for Comparing Categorical Maps. *International Journal of Remote Sensing*, 27(7):1425 – 1452, 2006.
[33] D. G. S. Richard O. Duda, Peter E. Hart. *Pattern Classification*. Wiley Interscience, 2 edition, 2000.
[34] G. Salamunićcar and S. Lončarić. GT-57633 catalogue of Martian impact craters developed for evaluation of crater detection algorithms. *Planetary and Space Science*, 56:1992–2008, Dec. 2008.
[35] G. Salamunićcar and S. Lončarić. Open framework for objective evaluation of crater detection algorithms with first test-field subsystem based on MOLA data. *Advances in Space Research*, 42:6–19, July 2008.
[36] M. Sharifzadeh, C. Shahabi, and C. Knoblock. Learning approximate thematic maps from labeled geospatial data. In *In the Proceedings of International Workshop on Next Generation Geospatial Information,Cambridge (Boston), Massachusetts, USA.*, 2003.
[37] J. Shi and J. Malik. Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8):888–905, 2000.
[38] D. Smith, G. Neumann, R. Arvidson, E. Guinness, and S. Slavney. Mars global surveyor laser altimeter mission experiment gridded data record. *NASA Planetary Data System*, (MGS-M-MOLA-MEGDR-L3-V1.0), 2003.
[39] T. Stepinski and R. Vilalta. Digital topography models for martian surfaces. *IEEE Geoscience and Remote Sensing Letters*, 2(3):260–264, July 2005.
[40] T. F. Stepinski, S. Ghosh, and R. Vilalta. Automatic recognition of landforms on mars using terrain segmentation and classification. In *Discovery Science*, pages 255–266, 2006.
[41] T. F. Stepinski, S. Ghosh, and R. Vilalta. Machine learning for automatic mapping of planetary surfaces. In *The Nineteenth Innovative Applications of Artificial Intelligence Conference*, 2007.
[42] K. L. Tanaka. The venus geographic mappers' handbook. *US Geol. Surv. Open File Rep.*, pages 99–438, 1994.
[43] J. Theiler. A contiguity-enhanced k-means clustering algorithm for unsupervised multispectral image segmentation. In *In Proceedings of SPIE 3159*, pages 108–118, 1997.
[44] R. Vilalta and T. F. Stepinski. Thematic maps of martian topography generated by a clustering algorithm. In S. Mackwell and E. Stansbery, editors, *Lunar and Planetary Institute Conference Abstracts*, pages 1169–1170, Mar. 2004.
[45] V. Walter. Object-based classification of remote sensing data for change detection. *International Journal of Photogrammetry and Remote Sensing*, 58:225–238, Jan. 2004.
[46] J.-P. Wang. Stochastic relaxation on partitions with connected components and its application to image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(6):619–636, 1998.
[47] P. G. Wetzler, R. Honda, B. Enke, W. J. Merline, C. R. Chapman, and M. C. Burl. Learning to detect small impact craters. In *WACV-MOTION '05: Proceedings of the Seventh IEEE Workshops on Application of Computer Vision (WACV/MOTION'05) - Volume 1*, pages 178–184, Washington, DC, USA, 2005. IEEE Computer Society.
[48] D. E. Wilhelms. Geologic mapping. *Planetary Mapping*, pages 208–260, 1990.
[49] I. H. Witten and E. Frank. *Data Mining: Practical Machine Learning Tools and Techniques, Second Edition (Morgan Kaufmann Series in Data Management Systems)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2005.

Soumya Ghosh is a graduate student pursuing PhD in the Department of Computer Science at the University of Colorado at Boulder. His research interests lie in exploring the intersection of machine learning and computer vision. He is interested in identification, recognition and characterization of objects in cluttered natural scenes and in combining these algorithms to achieve higher level artificial intelligence tasks such as Image Understanding. His current work focuses on developing machine learning solutions to problems in mobile robotics. His E-mail address is soumya.ghosh@colorado.edu

Tomasz Stepinski Tomasz Stepinski is a research staff member at the Lunar and Planetary Institute (Houston TX). He is a planetary scientist with a PhD degree in Applied Mathematics from the University of Arizona. His main area of research is automated identification, characterization, and classification of landforms, especially on the planet Mars. He leads the planetary science community in developing methods for automatic cataloging of impact craters and valley networks – prominent landforms on Mars. His E-mail address is tom@lpi.usra.edu

Ricardo Vilalta is an associate professor in the department of computer science at the University of Houston. He holds a MS and Ph.D. degrees in computer science from the University of Illinois at Urbana-Champaign. His research interests are in machine learning, statistical learning theory, data mining, and artificial intelligence. He is recipient of the best paper award at the European Conference on Machine Learning (2003) and recipient of the NSF Career Award (2005). His E-mail address is vilalta@cs.uh.edu