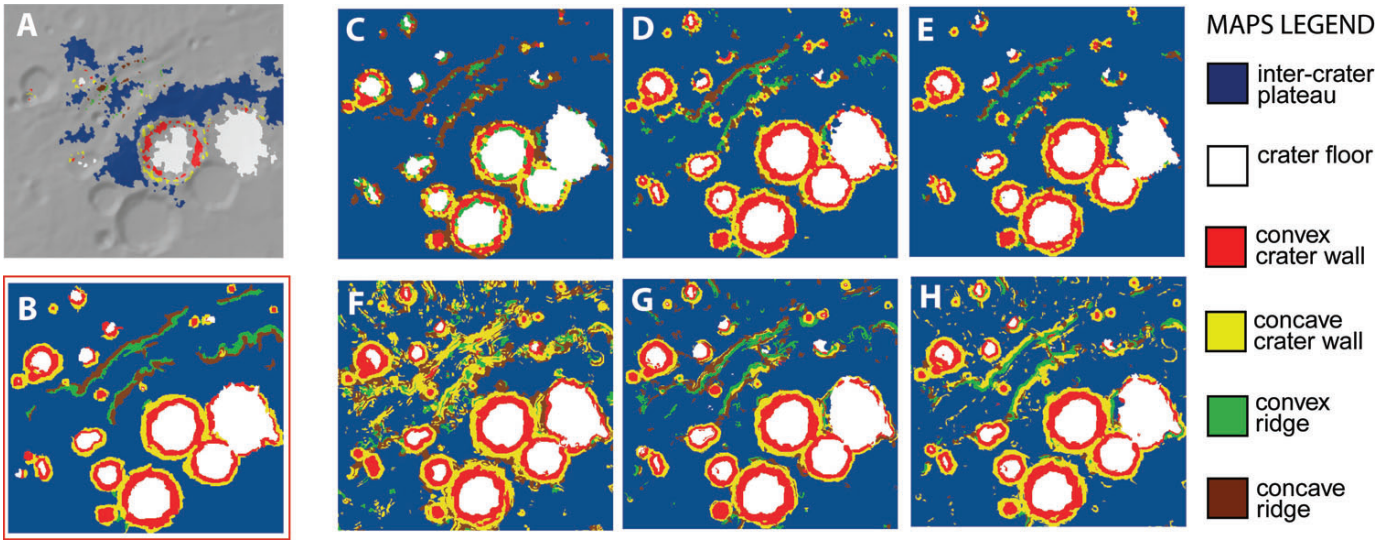

Fig. 4. Geomorphic maps of the Tisia site generated by different segmentation/classification methods. (A) Labeled regions; different landform classes are coded by different colors as given by the legend. Unlabeled regions are displayed in gray. (B) The analyst-drawn map ("ground truth"). Maps based on 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