

Predicting HiRISE-equivalent Rock Density on Mars Using CTX Image Features

Navid Serrano¹

Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, 91109

Patrick McGuire² and David Mayer³

Earth and Planetary Sciences, Washington University, St. Louis, MO, 63130

Andres Huertas⁴

Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, 91109

and

Raymond Arvidson⁵

Earth and Planetary Sciences, Washington University, St. Louis, MO, 63130

In this paper, a Bayesian framework is used to infer HiRISE-equivalent rock density from CTX image features without needing to explicitly identify rocks. The statistical relationship between HiRISE rock density and CTX image features is modeled using Bayesian Networks. The model is enriched by including geological features—specifically, geomorphic units. The geomorphic units are identified by geologists based on different surface features visible in a variety of orbital data. The inclusion of geomorphology in the Bayesian framework makes it possible to not only infer HiRISE-equivalent rock density directly from CTX image features, but also to make a preliminary assessment of the representative geomorphic units. HiRISE and CTX images of the Mars Phoenix mission landing area are used in this study. Rock density estimates from HiRISE images are used as ground truth to train the Bayesian Network using features extracted from corresponding CTX images. The approach is evaluated on a few test cases using a probabilistic estimate of landing safety based on inferred rock density from CTX image features. The initial results of this exploratory study show that predicted landing safety from the inferred rock density is very close to that obtained using actual HiRISE rock density estimates.

I. Introduction

Landing site selection is an essential part of the planning process for a space exploration mission. This process involves identifying regions of a planetary surface that meet the scientific objectives of the mission while ensuring the payload can be delivered safely. Identifying landing sites that meet the often stringent safety criteria for a mission is a painstaking process¹. Planetary orbiters deployed years before a potential surface mission are able to capture reconnaissance data that are indispensable for judiciously selecting a landing site. In addition to providing information on sites of scientific interest, orbital data can reveal key landing hazards that must be averted in order to ensure vehicle safety. Among potential landing hazards, rocks represent a significant risk. Abundant rocks at candidate landing sites were a dominant concern for the recent Mars Phoenix mission². Mission planners were able to quantitatively assess rock hazards at candidate landing sites with the aid of images from the High Resolution Imaging Science Experiment³ (HiRISE) onboard NASA's Mars Reconnaissance Orbiter (MRO). The unprecedented

¹ Member of Technical Staff, Computer Vision Group, M/S 198-219.

² Senior Research Fellow, McDonnell Center for the Space Sciences.

³ Research Assistant, Department of Earth and Planetary Sciences.

⁴ Senior Member of Technical Staff, Computer Vision Group, M/S 198-235.

⁵ James S. McDonnell Distinguished University Professor, Department of Earth and Planetary Sciences.

resolution (0.3 m/pixel) of HiRISE makes it possible to identify submeter-scale rocks that may pose a hazard to landing spacecraft. The ability to identify individual rocks in HiRISE images proved crucial in selecting a suitable landing site for the Phoenix mission. The development of a robust automatic rock detection algorithm⁴ provided a further boon for the Phoenix landing site selection process⁵. For the Phoenix mission, the rock detection algorithm was used to automatically count ~10,000,000 rocks from HiRISE images, which in turn were used to estimate the rock density within 100 m × 100 m regions of the Martian surface and ultimately predict the probability of landing success.

Despite their high resolution, HiRISE images have a smaller span relative to images from other existing cameras (e.g. Mars Orbiter Camera⁶) and thus cover less surface area. Although orbital data exists for areas of the Martian surface that HiRISE has not imaged, the resolution is much lower, making it impossible to resolve smaller rocks that still pose a significant threat to the safety of the spacecraft. Although individual rocks may not be visible in lower resolution images, there may be salient features—particularly texture signatures—that can be used to predict rock density. The purpose of this work is to explore a solution to this problem—namely, a means to predict HiRISE-equivalent rock density on the Martian surface from low resolution data. The low resolution data considered here is captured from the Context Imager⁷ (CTX), which has a lower resolution (6 m/pixel) than HiRISE images but has a broader span (24 km × 24 km). Moreover, an important characteristic of CTX and HiRISE images is that they are captured concurrently. Thus, there is a natural pairing between the two data sets.

It should be noted that the proposed approach is not necessarily intended to be trained on one area of Mars and tested in an entirely different area (unless the geomorphology and underlying rock formations are consistent). Rather, the goal would be to apply the approach in a region with scattered HiRISE coverage in order to “fill the gaps,” as shown in Fig. 1.

The problem is formulated using a Bayesian framework. Specifically, a Bayesian Network⁸ (BN) is used to graphically model the statistical relationship between rock density estimated from HiRISE images and features extracted from CTX images. In addition, geological indicators are incorporated into the BN in order to further enrich the description of the overall problem domain. In such a framework, the rock density can be inferred directly from CTX image features. Ultimately, the predicted rock density is used to determine the landing safety of a site. The images used in this study correspond to the Phoenix landing site⁵. The experimental results of this initial study demonstrate that the predicted landing safety from CTX image features is very close to the landing safety from HiRISE rock density for the same regions of Mars.

II. HiRISE and CTX Images

HiRISE is currently the highest resolution camera in operation around another planet. HiRISE is capable of imaging the Martian surface in near-true color using CCDs sensitive to red (570 nm-830 nm), blue-green (< 580 nm) and near infrared (NIR, >790 nm) light. The instrument has an optimal spatial resolution of 0.3 m/pixel from a 300 km orbit. The red CCDs have a field of view (FOV) of 1.14°, giving HiRISE images a maximum swath width of ~6 km. The along-track extent of HiRISE images is 18 km. A HiRISE color composite image has an FOV of 0.23°, corresponding to a swath width of 1.2 km³.

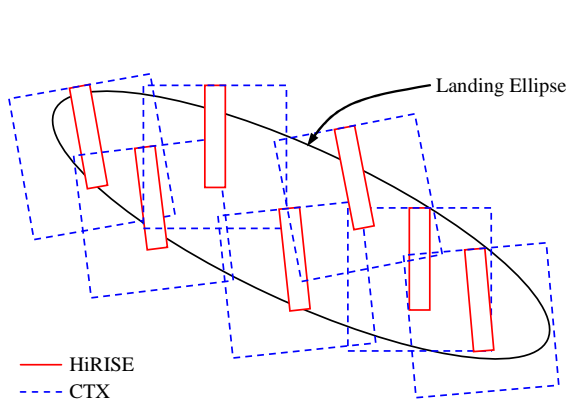


Figure 1. Example scenario with limited HiRISE and full CTX coverage within a landing ellipse.

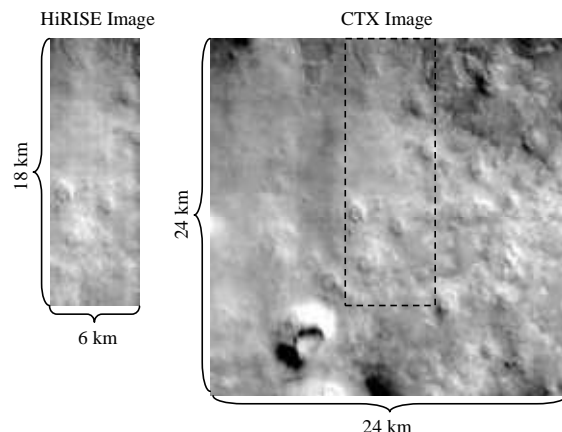


Figure 2. HiRISE and CTX image pair (PSP_1946) of the Phoenix landing area on Mars.

CTX is a single-channel, panchromatic camera sensitive to visible light within the 500 nm to 700 nm range. The camera uses a 5000 element CCD and is capable of imaging the Martian surface at an optimal resolution of 6 m/pixel from a 300 km orbit. CTX images have a nominal cross-track and along-track extent of 30 km⁷. The spatial resolution provided by CTX makes it especially useful for the study of surface rock distribution patterns at the scale of several meters to tens of meters.

HiRISE and CTX are capable of acquiring images concurrently, thus there is a natural pairing between the two datasets. This makes co-registering images from HiRISE and CTX for the purpose of extrapolating rock distribution simpler and more accurate than attempting to co-register HiRISE with images from instruments with higher resolutions than CTX (such as meter-scale images from the defunct Mars Orbiter Camera). An illustrative example is shown in Fig. 1 with scattered HiRISE coverage within a landing ellipse. As can be seen from Fig. 1, the CTX images cover the entire landing ellipse and therefore could be used to predict the rock density for regions not covered by HiRISE. An example HiRISE and CTX image pair is shown in Fig. 2.

III. Geologic Setting of Study Area

A. Regional Geology

The Phoenix landing site was selected for this study due to the availability of HiRISE and CTX image pairs and the familiarity of the authors with the area. The area is located in the Borealis basin in the northern plains of Mars. Regional slopes are generally less than 0.1° and dip towards the north⁹. The landing area is situated over two distinct geologic units within the Vastitas Borealis Formation: the Vastitas Borealis marginal unit and Scandia region unit¹⁰. The Vastitas Borealis marginal unit forms smooth plains with occasionally dissected low plateaus. This unit is interpreted as a sedimentary deposit composed of material transported through a system of outflow channels originating in the southern highlands⁹. The Scandia region unit overlies the Vastitas Borealis marginal unit and is characterized by knobby terrain, interspersed with irregular topographic depressions⁹. This unit is thought to have formed by a variety of subsurface volatile processes, possibly driven by geothermal heating from the nearby Alba Patera volcano¹⁰.

B. Geomorphology of the Phoenix Landing Area

Seelos et al.⁹ produced a series of geomorphic maps of candidate landing sites for the Phoenix mission as part of the landing site selection process. The map around the landing site that was ultimately chosen extends from 67°N to 68.5°N latitude and 229°E to 238°E longitude (Fig. 3). The map was produced using visible images from the THEMIS instrument¹¹ combined with MOLA topography¹² as basemaps. Geomorphic units were described on the basis of their topographic expression and relative albedo differences. Although the THEMIS and MOLA basemaps provided complete coverage of the landing area, their low resolutions (18 m/pixel to 36 m/pixel for THEMIS and ~231 m/pixel for MOLA) relative to HiRISE and CTX limited the resolvability of features in the geomorphic map to 50 m. HiRISE and CTX coverage remains limited within the map in Fig. 3.

Seelos et al.⁹ defined seven geomorphic units at the 50 m scale within the Phoenix landing area: Highlands, blocks/mesas, knobs, lowlands bright, lowlands dark, crater interiors and crater ejecta. The highlands unit is characterized by smooth intercrater plains located at a distinctly higher topographic position than the lowlands units (Fig. 3). The blocks/mesas unit is represented by flat-topped, elevated landforms with relief of a few meters relative to the surrounding terrain. Individual mesas range from ~5 km to 50 km in diameter. The knobs unit is represented by rounded, often clustered hills with relief of several tens of meters to several hundred meters. Individual knobs have basal diameters on the order of a few kilometers. The lowlands unit is divided into light and dark subunits distinguished by their albedo differences. Both subunits consist of shallow, irregularly-shaped depressions with widths of several tens of kilometers. Impact craters and related ejecta deposits are distributed across the entire Phoenix landing area. The crater interior geomorphic unit is identified by the presence of bowl-shaped or circular impact structures. The crater ejecta unit is composed of material emplaced by impact processes and forms radial to braided distribution patterns or circumferential ramparts outside craters.

C. Importance of Rock Distribution for Landing Site Selection

The assessment of candidate landing sites for planetary surface missions requires mission planners to balance potential science return against landing site hazards. Rocks on the surface of a landing site can damage a spacecraft during landing or pose navigational hazards to mobile surface missions. The hazards posed by surface rocks necessitate the development of rock distribution maps for candidate landing sites as part of the landing site assessment process.

Engineering constraints for the Phoenix mission stipulated that a landing site must be free of rocks down to 0.33 m tall in order to allow the spacecraft's solar panels to fully deploy and avoid puncturing the bottom of the spacecraft during landing⁵. Extensive submeter-scale imaging by HiRISE was required in order to explicitly identify rocks within candidate landing areas for the mission. Rock distribution maps based on these HiRISE images were used to select the final landing site, resulting in a successful landing of the spacecraft.

However, future spacecraft may be unable to depend upon submeter-scale images of candidate landing sites due to the logistical difficulty in covering large areas at high resolution. This possibility requires the development of alternative methods for quantifying rock distribution at candidate landing sites.




Figure 3. Geomorphic map of Phoenix landing area. *Geomorphic units are highlands (H), blocks/mesas (B/M), knobs (K), lowland dark (Ld), lowland bright (Lb), crater ejecta (Ce), crater interior (Ci).* (Modified from Seelos et al.⁴.)

IV. Estimating Rock Density

A. Traditional Approaches

Considerable work has been done over the years to model the distribution of rocks on Mars¹³. For instance, size-frequency distributions of rocks have been derived based on images of the Martian surface captured by the Viking Landers and compared against analogous sites on Earth¹⁴. For landing site selection, rock abundance predicted from remote sensing data has been compared against derived size-frequency distribution models for validation. In particular, thermal inertia measurements obtained from instruments like THEMIS have been used to predict surface rock abundance for landing site selection, as in the case of the Mars Exploration Rovers mission¹. However, during the landing site selection process for the Phoenix mission, the first HiRISE images revealed far higher rock abundances than previously believed². Hence, characterization of rock abundance directly from HiRISE images emerged as an area of critical need in order to identify a suitable landing site for the Phoenix mission. Initial rock counts were performed by hand² but the sheer volume of rocks in the area would have made it impossible to complete the task prior to the target launch date. Fortunately, a robust automatic rock detection algorithm (see below) was used to identify individual rocks in HiRISE images and as a result, rock density maps were produced for most of the area within the Phoenix landing ellipse⁵.

B. Automatic Rock Detection from HiRISE Images

The automatic rock detection algorithm used to estimate rock density in the Phoenix landing area was first envisioned for on-board processing during spacecraft entry, descent, and landing⁴ but was shown to be highly effective at identifying rocks in HiRISE images⁵. The technique identifies rocks in two steps: 1) shadow extraction, and 2) shadow analysis.

Shadow extraction involves grouping dark regions of the image using a technique based on Maximum Entropy Thresholding with gamma correction (gMET). Thresholding is a simple way to segment an image into different sets based on pixel intensity. Using the image histogram as a representation of the distribution of pixel intensities, the gMET algorithm chooses the threshold that maximizes the entropy between shadow and nonshadow regions. The histogram of pixel intensities is often not grouped into discernable shadow and nonshadow regions; therefore, gamma correction is applied to the HiRISE image in order to force bimodality.

Shadow analysis is the key step for rock modeling. The rock-shadow model is shown in Fig. 4. As can be seen from Fig. 4, the extracted rock shadow is fit with an ellipse that preserves the shadow area. The semi-major axis of the ellipse (dashed line) closely follows the direction of the illumination ray (solid line) and thus its length provides an estimate of the shadow length. The rock itself is modeled as a cylinder with a radius b equal to the semi-minor axis length of the shadow ellipse and centered at the terminator. The cumulative fractional area of rock coverage on Mars is typically represented as a function of rock diameter. Hence, the estimated diameter of a detected rock is simply $D=2b$. The shadow length and known sun incidence angle are used to estimate rock height.




Figure 4. Rock model used to estimate rock size from extracted shadow region.

For the Phoenix mission, the rock detection algorithm was used to identify rocks with a diameter greater than 1.5 m in HiRISE images. The rock detection results are reported in terms of diameter and (x,y) locations in the HiRISE image. With these results, the rock density can be estimated by simply counting the number of rocks within a region of a given spatial extent. For the Phoenix mission, the rock density was computed over $100\text{ m} \times 100\text{ m}$ (hectare) regions in the HiRISE images. An example HiRISE image and its corresponding rock density map obtained from automatic rock detection are shown in Fig. 5.

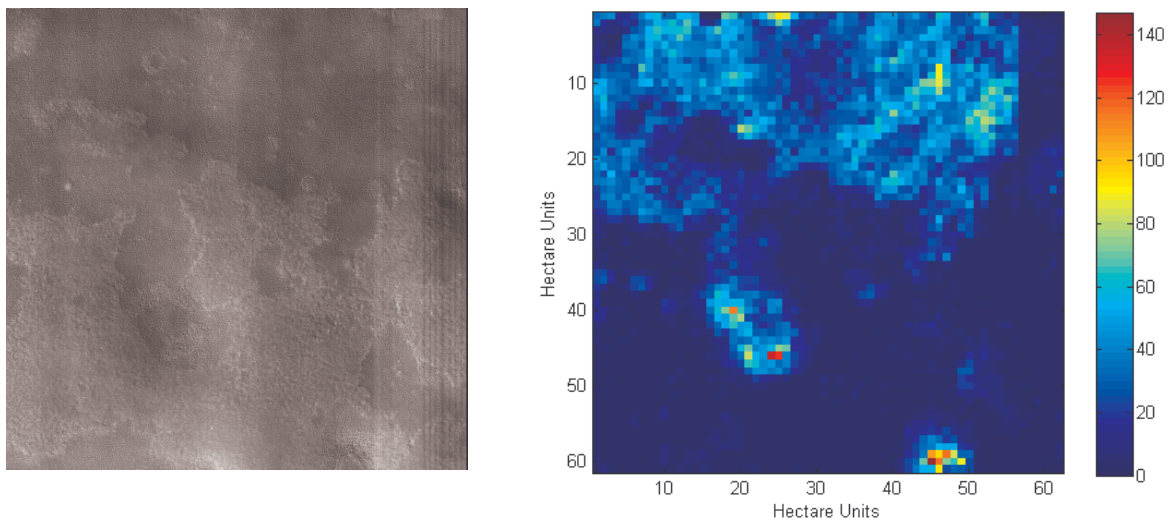


Figure 5. Rock density from HiRISE image. Original HiRISE image (left) and corresponding rock density—number of rocks larger than 1.5 m per hectare (right).

V. Predicting Rock Density from CTX Images

A. Feature Extraction

Although individual rocks are not visible in CTX images, the underlying rock formations manifest themselves as discernable textures and intensity variations. These patterns are also consistent with the geomorphic units defined in Fig. 3. Hence, intuitively, it should be possible to extract features from CTX images that correlate well with the underlying rock density.

The texture in CTX images is modeled using Gray Level Co-occurrence Matrix (GLCM) features¹⁵. Let P be an $L \times L$ matrix. Each element p_{ij} in P represents the number of times gray levels i and j occur with a displacement (dx, dy) :

$$p_{ij} = \sum_m \sum_n \begin{cases} 1, & \text{if } I(m, n) = i \text{ and } I(m + dx, n + dy) = j \\ 0, & \text{otherwise} \end{cases}, \quad (1)$$

where $I(x, y)$ is the gray level (CTX) image. In this case, GLCM features are computed at an orientation of 135° ; that is, $dx=1$ and $dy=1$. The size of the matrix P depends on the number of gray levels L in the image $I(x, y)$. In order to keep P reasonably small, the GLCM features are computed for $L=16$ levels. Given a co-occurrence matrix P , a variety of statistical features can be computed. For this study, four GLCM texture features were considered:

$$f_1 = \sum_i \sum_j (i - j)^2 p_{ij}, \quad (2)$$

$$f_2 = \sum_i \sum_j p_{ij}^2, \quad (3)$$

$$f_3 = \sum_i \sum_j \frac{p_{ij}}{1 + |i - j|}, \quad (4)$$

$$f_4 = \frac{1}{\sigma_m \sigma_n} \sum_i \sum_j (i - \mu_m)(i - \mu_n) p_{ij}, \quad (5)$$

where f_1, \dots, f_4 represent the contrast, energy, homogeneity, and correlation of P , respectively, and $\mu_m, \mu_n, \sigma_m,$ and σ_n are the marginal means and standard deviations along the rows and columns of P . The GLCM features are extracted over an observation window of size $W=16$, which given the 6 m/pixel resolution of CTX images, corresponds to a $100 \text{ m} \times 100 \text{ m}$ (or hectare) sized region. This ensures that the CTX features are extracted from the same spatial region as the HiRISE rock density. Simple image statistics are also computed in order to capture local intensity variations. Hence, two additional features $f_5 = \mu_I$ and $f_6 = \sigma_I$ are used as predictors, where μ_I and σ_I are the mean and standard deviation of the CTX image $I(x, y)$ over the observation window W . An example feature set extracted from a CTX image is shown in Fig. 6. In addition, the rock density computed from the corresponding HiRISE image is also shown. As can be seen, the CTX image features visually correlate with the HiRISE rock density computed over the same region of Mars.

B. Bayesian Formulation of the Inference Problem

A Bayesian network (BN) is a directed acyclic graph (DAG)⁸ in which the nodes represent variables and the links between nodes represent causal dependence expressed by conditional probability distributions (CPDs). A link originates at a parent node and is directed toward a child node. The direction of the link indicates causality, and thus a dependence relationship. Nodes that exist at the same level are considered conditionally independent. Such a framework can be regarded as a knowledge representation because it encodes the joint probability of the variables. An important feature of BNs is that computation of the joint probability is simplified by taking advantage of the conditional independencies encoded in the graph.

The structure of a BN is usually derived based on domain knowledge of the relationships between different variables. The Bayesian formulation proposed here does not merely consider the relationship between HiRISE rock

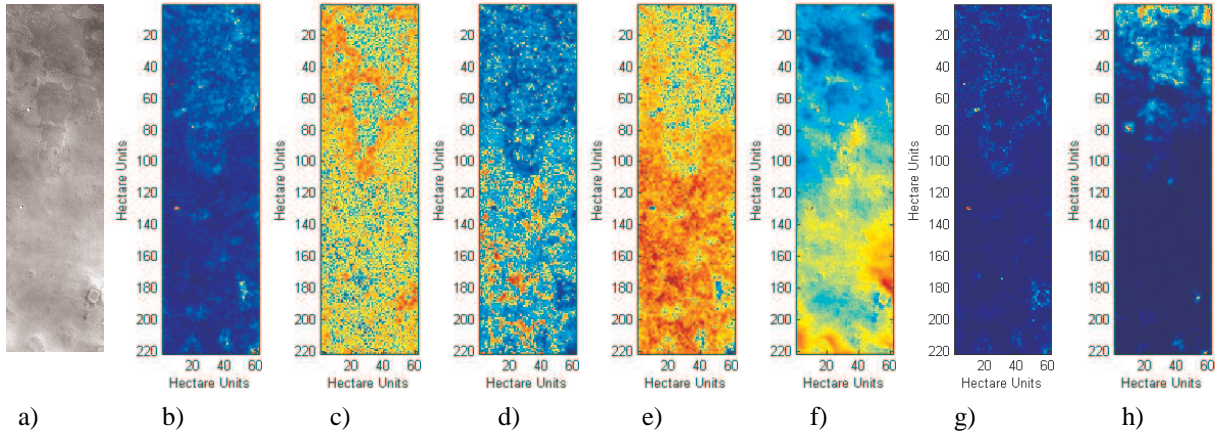


Figure 6. Example CTX image features and corresponding HiRISE rock density. *a) original image, b) contrast, c) correlation, d) energy, e) homogeneity, f) mean, g) standard deviation, and h) HiRISE rock density.*

density and CTX image features but also incorporates geomorphology. As discussed earlier, the geomorphic units defined for the Phoenix landing area are related to the underlying rock formations. Hence, it is useful to include geomorphology in the Bayesian framework. Not only will this reveal information on the statistical relationship between rock density and the geomorphic units but it will also make it possible to infer the geomorphic unit of a region from CTX image features directly. This could potentially be very useful since geomorphic maps (as in Fig. 3) are an important tool for geologists but typically involve considerable work to produce. The ability to infer a probable geomorphic unit from CTX image features could serve as a preliminary step in the overall process of defining a geomorphic map.

Let D , G , and f_1, \dots, f_N be random variables representing rock density, geomorphic unit, and CTX image features, respectively. (In this case, the number of features is $N=6$, as discussed in the previous section.) The relationship between these variables is shown graphically in Fig. 7. A BN encodes the joint probability distribution between the variables and reveals statistical independencies. Hence, the joint probability for the BN in Fig. 7 can be written as:




Figure 7. BN graphical structure used to model the relationship between HiRISE rock density, D , geomorphic unit, G , and CTX image features f_1, \dots, f_N . *The nodes that are hidden during inference are shaded.*

$$P(D, G, f_1, \dots, f_N) = P(D)P(G | D) \prod_{i=1}^N P(f_i | G), \quad (6)$$

As can be seen from Eq. (6), the CTX image features f_1, \dots, f_N are assumed to be conditionally independent. This assumption has generally been shown to be adequate for similar problems and does not lead to increased error¹⁶.

Inference in a BN involves computing the probability at a particular node given a set of evidence. In this case, there are two variables of interest: rock density D and geomorphic unit G . Only the CTX image features f_1, \dots, f_N are observable during inference, and thus, rock density D and geomorphic unit G are treated as hidden nodes, as shown in Fig. 7. Having learned the CPDs between the variables in the BN (see following section), the probability at nodes D and G can be inferred directly from the features by computing $P(D | f_1, \dots, f_N)$ and $P(G | f_1, \dots, f_N)$, respectively. These two posterior probabilities are the quantities of interest and represent the probability of a particular rock density given observed CTX image features and the probability of a particular geomorphic unit given observed CTX image features, respectively.

C. Statistical Learning

Statistical learning involves computing the CPDs from a set of training observations. During the training process all nodes are observable. That is, in addition to the CTX image features f_1, \dots, f_N , the BN is also populated with HiRISE rock density ground truth and the known geomorphic unit obtained from the map in Fig. 3. The HiRISE rock density D and the CTX image features f_1, \dots, f_N are all continuous-valued variables with unknown distributions. Instead of assuming a particular continuous distribution (e.g. Gaussian), all continuous variables are discretized and instead the CPDs are represented as conditional probability tables (CPTs). This has several advantages; one, it approximates the arbitrary shape of the CPDs, and two, it makes training and inference in the BN much more efficient.

The optimal choice of the discretization depends on the available data and the particular problem. In this case, an information theoretic approach is used to determine the optimal number of discretization levels; specifically, the Bayesian Information Criterion¹⁷ (BIC) is used to estimate the dimension of each continuous variable. Training is carried out using the Bayes Net Toolbox¹⁸ (BNT).

The HiRISE and CTX images used in this study are listed in Table 1. All of these images are from the Phoenix landing area shown in Fig. 3. As discussed earlier, the HiRISE and CTX images are captured concurrently resulting in a natural pairing. However, because they do not cover the same span, they need to be registered. For registration purposes, the HiRISE images were downsampled to 6 m/pixel (the CTX image resolution) and registered based on prominent features visible in both images (e.g. craters). The CTX images were then cropped so that they overlap completely with the corresponding HiRISE images. This was done in order to ensure a completely observable data set for training. In practice, the CTX images need not be cropped once the BN is trained; the rock density can be estimated from the entire image, which is, ultimately, the main objective of this work.

Typically, large training sets are required in order to achieve generalization. Unfortunately, the number of HiRISE-CTX image pairs available for this study is relatively small. Hence, more cases (80%) were included in the training set for better generalization. As shown in Table 1, training was performed on 8 of the 10 images and the remaining 2 were used for testing. Most of the HiRISE and CTX images used in this study are composed of lowland bright (Lb) and lowland dark (Ld) material (see Table 1). There are also some prominent craters (Ci) and related ejecta (Ce). One image (PSP_1893) also contained knobby (K) terrain. However, it was excluded from the set used in this study because there would have been too few representative cases of this particular type of terrain.

Table 1. HiRISE-CTX image pairs used in the study.

Case	HiRISE-CTX Image Pair	Geomorphic Units Present
Training	PSP_1880	Lb, Ld, Ci, Ce
	PSP_1906	Lb
	PSP_1959	Lb, Ld, Ci, Ce
	PSP_1972	Lb, Ld, Ci, Ce
	PSP_2104	Lb, Ci, Ce
	PSP_2170	Lb, Ld, Ci, Ce
	PSP_2183	Lb
	PSP_2249	Lb
Testing	PSP_1946	Lb, Ld, Ci, Ce
	PSP_2012	Lb, Ld, Ci, Ce

D. Expected Landing Safety from Rock Density

Once the CPTs in the BN are learned, inference can be performed on unobserved test cases. The rock density can be inferred from CTX image features using the posterior probability $P(D | f_1, \dots, f_N)$. In practice, candidate sites are selected based on landing safety, which is measured by the probability of encountering various hazards. Often, hazard maps are used to visualize landing safety. For example, Table 2 shows the hazard map color coding based on rock density that was used during the Phoenix landing site selection process. The acceptable to unacceptable rock density ranges shown in Table 2 are used to compute a hazard map from CTX image features that is analogous to the hazard maps produced from HiRISE rock density. The expected landing safety from CTX image features is:

$$E(S | f_1, \dots, f_N) = \sum_i s_i P(D \in R_i | f_1, \dots, f_N), \quad (7)$$

where, S is landing safety, s_i is the landing safety value corresponding to rock density range R_i , Based on Table 2, the rock density ranges are defined as $R_1=[0,3]$, $R_2=(3,8]$, $R_3=(8,19]$, and $R_4=(19,\infty)$. The expected landing safety $E(S|f_1,\dots,f_N)$ from CTX image features can be seen as a continuous-valued hazard map and is used to compare against the hazard maps obtained from HiRISE rock density using Table 2.

A similar approach can be used to predict the geomorphic unit of a region (e.g. hectare) based on CTX image features. The number of geomorphic units is finite and thus the most likely geomorphic unit can be obtained from the posterior $P(G|f_1,\dots,f_N)$:

$$\hat{g} = \arg \max_g P(G = g | f_1, \dots, f_N), \quad (8)$$

Although the actual geomorphic unit for a region may be predicted using Eq. (8), it may be more useful to define a measure that captures uncertainty about the geomorphology. This measure can then be used by scientists to make a final determination of the actual geomorphic unit. As with landing safety, an expectation can be used:

$$E(G | f_1, \dots, f_N) = \sum_g g P(G = g | f_1, \dots, f_N), \quad (9)$$

Table 2. Hazard map color coding based on rock density.

Color	Rock Density Per Hectare
Green	0 – 3
Yellow	4 – 8
Orange	9 – 19
Red	> 19

E. Experimental Results

Two CTX images were excluded from the training set and are used to test the proposed approach. The posterior probability $P(D|f_1,\dots,f_N)$ is inferred for each hectare in an image by propagating the feature values f_1,\dots,f_N through the BN. The expected landing safety can then be computed with Eq. (7). The original CTX image, CTX hazard map, HiRISE hazard map, and expected geomorphic unit are shown for test cases PSP_2012 and PSP_1946 in Fig. 8 and Fig. 9, respectively.

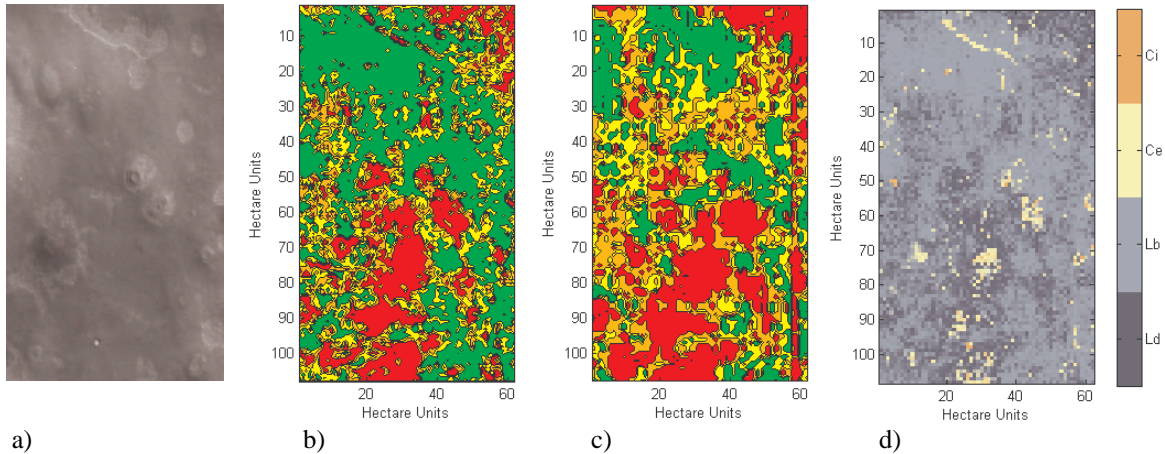


Figure 8. Results for PSP_2012. a) original CTX image, b) hazard map from CTX image features, c) hazard map from HiRISE rock density, and d) expected geomorphic unit per hectare from CTX image features.

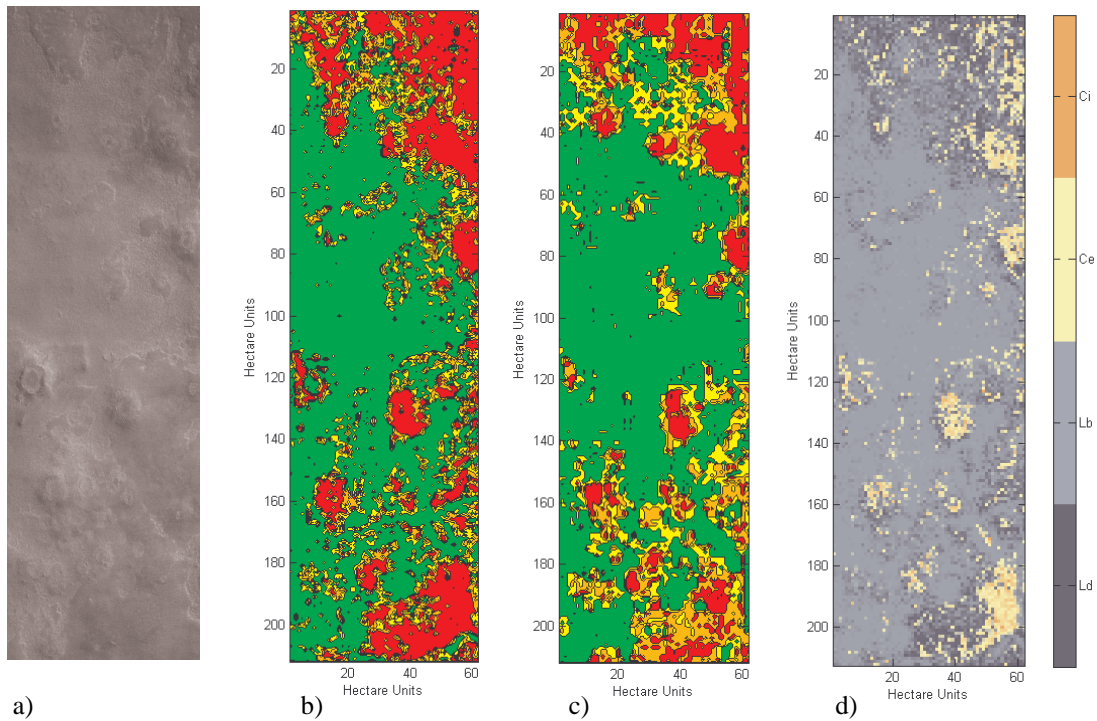


Figure 9. Results for PSP_1946. *a) original CTX image, b) hazard map from CTX image features, c) hazard map from HiRISE rock density, and d) expected geomorphic unit per hectare from CTX image features.*

As can be seen from the results in Fig. 8-9, the hazard map from CTX image features is very close to the hazard map computed from HiRISE rock density for both test cases. For PSP_2012 (Fig. 8), the CTX hazard map tends to under-estimate the underlying rock density in a few regions. This is most likely because the BN identifies some Ld regions as Lb (Fig. 8d), which leads to a lower rock density prediction in a few small regions. For PSP_1946 (Fig. 9), there are some regions of the CTX hazard map that tend to over-estimate the underlying rock density compared to the HiRISE hazard map. In this case, the BN identifies some Ce areas that are actually Ld (Fig. 9d). Since the rock density around crater ejecta deposits tends to be high; the BN accordingly predicts a higher rock density. This can be seen in the upper right and lower right portions of Fig. 9d. Despite these minor details, it can be argued that the CTX hazard maps are remarkably close to the HiRISE hazard maps considering they are inferred from lower resolution image features.

It should be noted that the expected geomorphic unit shown in Fig. 8d and Fig. 9d is on a continuous scale. In other words, the actual geomorphic unit is not being predicted. Probabilistically, the map can be interpreted as the most likely geomorphic unit given a set of CTX image features. The maps in Fig. 8d and Fig. 9d could serve as a preliminary assessment of the geomorphology for regions that have not been mapped manually by scientists. Considering how demanding it is to produce a geomorphic map, the ability to quickly obtain a preliminary assessment might be very attractive. It should further be noted that all of these results could be improved with a richer dataset.

VI. Conclusion and Future Work

This paper presented an initial exploration of the problem of inferring rock density from CTX image features at a level close to or equal to that obtained from automatic rock counts in HiRISE images. A Bayesian Network was used to model the relationship between HiRISE rock density, geomorphology, and CTX image features. The resulting probabilistic framework allows for computation of CTX hazard maps that are analogous to the hazard maps produced using rock density estimates from HiRISE images. The initial results are promising and show that the CTX and HiRISE hazard maps are visually correlated. The results suggest that such an approach, after further refinement, could serve to produce hazard maps from CTX images in areas that have not been imaged by HiRISE. Further work is needed in order to make the proposed approach more robust. First, a larger database would be

instrumental in generalizing the results beyond a few test cases. Second, other areas beyond the Phoenix landing area should be evaluated since the geology (and consequently, the visible features) would be quite different. In addition, it would be worthwhile to evaluate other CTX image features beyond those proposed here and also compare the Bayesian approach with other machine learning techniques.

Acknowledgments

The research described in this paper was carried out at the Jet Propulsion Laboratory (JPL), California Institute of Technology, under contract with the National Aeronautics and Space Administration (NASA), and the Earth and Planetary Remote Sensing Laboratory (EPRSL), Washington University in Saint Louis. Funding was provided by the JPL Director's Research and Development Fund. The authors would like to thank Yang Cheng and Shane Brennan of JPL for helpful discussions and comments and in particular, Doug Adams of JPL who provided much of the original impetus for this study.

References

- ¹Golombek, M.P., et al., "Selection of the Mars Exploration Rovers Landing Site," *Journal of Geophysical Research*, Vol. 108(E12), 8072, doi:10.1029/2003JE002074, 2003.
- ²Arvidson, R., et al., "Mars Exploration Program 2007 Phoenix Landing Site Selection and Characteristics," *Journal of Geophysical Research*, Vol. 113, E00A03, doi:10.1029/2007JE003021, 2008.
- ³McEwen, A. S., et al. "Mars Reconnaissance Orbiter's High Resolution Imaging Science Experiment (HiRISE)," *Journal of Geophysical Research*, Vol. 112, E05S02, doi:10.1029/2005JE002605, 2007.
- ⁴Huertas, A., Cheng, Y., and Madison, R., "Passive Imaging-Based Multi-Cue Hazard Detection for Spacecraft Landing," *IEEE Aerospace Conference*, Big Sky, MT, March, 2006.
- ⁵Golombek, M.P. et al., "Size-Frequency Distributions of Rocks on the Northern Plains of Mars with Special Reference to Phoenix Landing Surfaces," *Journal of Geophysical Research*, Vol. 113, E00A09, doi:10.1029/2007JE003065, 2008
- ⁶Malin, M.C., and Edgett, K.S., "Mars Global Surveyor Mars Orbiter Camera: Interplanetary cruise through primary mission," *Journal of Geophysical Research*, Vol. 106, E10, 23,429–23,570, 2001.
- ⁷Malin, M.C., et al., "Context Camera Investigation On Board the Mars Reconnaissance Orbiter," *Journal of Geophysical Research*, Vol. 112, E05S04, doi:10.1029/2006JE002808, 2007.
- ⁸Jensen, F.V., and Nielsen, T.D., *Bayesian Networks and Decision Graphs*, 2nd ed., Springer-Verlag, New York, 2007.
- ⁹Seelos, K.D., et al., "Geomorphologic and Mineralogic Characterization of the Northern Plains of Mars at the Phoenix Mission Candidate Landing Sites," *Journal of Geophysical Research*, Vol. 113, E00A13, doi:10.1029/2008JE003088, 2008.
- ¹⁰Tanaka, K.L., Skinner, A., and Hare, T.M., *Geologic Map of the Northern Plains of Mars*, *Science Investigation*, Map 2888, United States Geological Survey, Reston, VA.
- ¹¹Christensen, P.R., "The Thermal Emission Imaging System (THEMIS) for the Mars Odyssey 2001 Mission," *Space Science Reviews*, Vol. 110, 85-130, doi:10.1023/B:SPAC.0000021008.16305.94, 2004.
- ¹²Zuber, M., et al., "The Mars Observer Laser Altimeter Investigation," *Journal of Geophysical Research*, Vol. 97, 7781-7797, 1992.
- ¹³Christensen, P. R., "The spatial distribution of rocks on Mars," *Icarus*, Vol. 68, 217–238, 1986.
- ¹⁴Golombek, M., and Rapp, D., "Size-Frequency Distributions of Rocks on Mars and Earth Analog Sites: Implications for Future Landed Missions," *Journal of Geophysical Research*, Vol. 102, 4117– 4129, 1997.
- ¹⁵Haralick, R. M., Shanmugam, K., and Dinstein, I., "Textural Features for Image Classification," *IEEE Transactions on Systems, Man, and Cybernetics*, Vol. 3, No. 6, 610–621, 1973.
- ¹⁶Friedman, N., Geiger, D., and Goldszmidt, M., "Bayesian Network Classifiers," *Machine Learning*, Vol. 29, 131 – 163, 1997.
- ¹⁷Schwarz, G., "Estimating the Dimension of a Model," *Annals of Statistics*, Vol. 6, 461 – 464, 1978.
- ¹⁸Murphy, K., "The Bayes Net Toolbox for Matlab," *Computing Science and Statistics*, Vol. 33, 2001.