Tina Memo No. 2011-005
Internal

Towards a Quantitative Analysis of Martian Terrains

P.Tar and N.A.Thacker

Last updated
20 / 9 / 2011

Imaging Science and Biomedical Engineering Division,
Medical School, University of Manchester,
Stopford Building, Oxford Road,
Manchester, M13 9PT.

Towards a Quantitative Analysis of Martian Terrains

Paul Tar, Neil Thacker

Abstract

This document investigates the use of sample based Bayesian segmentation algorithms on images of the Martian landscape, with regard to a quantitative understanding of statistical and systematic errors. The errors predicted from stochastic sampling are estimated using the conventional methods of error propagation and the Cramer-Rao bound and compared to the estimation errors for computed total quantities of textures. We show that systematic errors will rapidy grow to exceed the statistical errors for cases where the additive texture model does not discribe the data density distribution accurately. We investigate how these distributions might be better modelled using mixtures. We conclude that more work is still needed to apply such pattern recognition methods to scientific measurement.

1 Introduction

Imaging of planetary surfaces is a high profile element of missions to explore the solar system. Over the last 10-15 years a series of orbiting missions have returned images of the Martian surface with unprecedented resolution (down to 0.25 m per pixel [1]), while the current missions Messenger [2] and Dawn [3] will provide high resolution coverage of the surfaces of Mercury and Vesta respectively. Such missions present an unrivalled opportunity to understand the history of geological processes that have shaped the surfaces of the terrestrial planets and ice moons.

Planetary surfaces respond to and record processes driven from both above and below the surface. The nature and relative importance of these processes varies from place to place and, in any one place, through time; the current morphology of the landscape is thus a product of (and hence record of) ancient and recent variations in processes, in some cases driven by climate change. For example, variation in the rate of uplift of a surface alters the slope of the land, affecting the direction and pattern of drainage networks, and degree of dissection of the surface. In geomorphological studies on Earth, the scale and pattern metrics of landscape organization are qualitatively and quantitatively analysed to determine the principal active processes responsible for the evolution of the terrain [4]. Recent developments in planetary imaging allow such techniques to be adapted to understand the evolution of other planets.

Automated image analysis has the potential to be a key component in this enterprise, especially in the case of Mars. The quantity of available imagery has surpassed the capacity of individuals to make first pass assessments seeking the presence or absence of textures diagnostic of particular processes. In addition to selecting promising candidates for further study, possible applications of an automated 'sifting' system include the identification of images containing evidence of drainage networks or dessication, whose distribution across a planetary surface is diagnostic of past or present climate patterns [5] [6]. There is also interest in the mapping of dunes and their relationship with climate [7]. Another potential application relates to the determination of relative surface ages from size frequency distributions (SFDs) of impact craters [8]. The ability to place geomorphic features in a chronological sequence, and hence to establish relative timing in landscape, is vital to developing and testing, quantitatively, models of landscape evolution.

For any data to be scientifically useful, outputs must be quantitatively valid, which for the purposes of this paper will be defined as 'reflecting underlying ground truths within clearly stated confidence intervals'. For example, summary outputs may involve plots showing the quantities of features present accompanied by error bars allowing results to be meaningfully compared to theory by scientific researchers. Automatically generated SFDs must be capable of generating such error bars to fulfil scientific researchers' needs. Achieving this quantitative validity for a variety of outputs, including SFDs, is the focus of our current research. On some level each of these applications involves the recognition and counting of pixels belonging to different classes of geological/geomorphological feature.

As a proof of concept this paper will focus on a simpler problem of pixel-wise analysis of Martian terrains at a fixed scale and orientation. The aim is to produce quantitative estimates of the amounts, in pixels, of different terrain types within a set of test images. This is achieved using a BRIEF-like local image descriptor to encode image patch information, a classifier based upon Bayes Theorem and Expectation Maximisation, and the application of a lower variance bound along with error propagation for the construction of error bars.

Existing work on the analysis of planetary surfaces has concenrated on the detection of impact craters using Hough transforms [10] [12] [13] [15] [16] [17], ellipse or circle fitting [14] [16], various types of template matching [10] [15], binary image processing on shadows [11], the Haar transform [18] [19], grey level co-occurrences and appearance

modelling [20]. To improve accuracy some methods pre-process images to determine regions of interest using grey level thresholds [11] and Grey Level Co-occurrence Matrices (GLCM) [14]. Some methods use post-processing such as morphological operations to join edges or remove isolated edges [9]. Despite such filtering many false positives are reported which have been further filtered using size thresholds [11], neural networks and models in eigenspace [14], or cross-referencing against alternative detection methods [15].

The segmentation of Martian terrains can be seen as a texture analysis problem. Recent BMVC work on texture analysis includes the use of Texture Fragmentation and Reconstruction (TFR) [24] and Markov Random Fields (MRF) [25] where textures are described in graph form. Textures have also been modelled using extended Fields-of-Experts [26] with performance demonstrated on synthesis tasks. Classification of textures has been performed using Random Forests (RF) by randomly sampled pixels [27], and Naive Bayes Classifiers (NBC) compared with Naive Credal Classifiers (NCC) using Local Binary Pattern descriptors [28]. In an attempt to mitigate against erroneous classifications the latter argues that returning a set of potential classifications from an NCC is better than forcing a single class decision.

Methods have also been used to extract other features from Digital Elevation Maps (DEM). These include river channels using paths of steepest decent [21] [22] and fault scarps using wavelets [23].

From a quantitative perspective, we are not aware of any reported methods which attempt to estimate quantitation error. Performance is reported only in absolute terms based upon discrete counts of successes and failures. A quantitative error assessment of craters detected using Support Vector Machines (SVM) has been made by comparing results against confidence intervals set by human performance [9]. However, this work did not explicitly generate confidence intervals for the SVM results.

2 Methodology

An 'off-the-shelf' solution to estimating the quantities of pixels within an image associated with different Martian textures could involve selecting a common representation (Grey Level Co-occurrence Matrix, Gabor filters etc.) and learning textures with a standard classifier (Support Vector Machine, Random Forest etc.). The output of the classifier could give a best-guess class label for each pixel which could then be counted. A Monte-Carlo simulation could be used to calibrate results to give approximate confidence intervals within which estimated quantities could be trusted, i.e. sample an empirical estimate of the covariance matrix. We propose an alternative method capable of correcting for misclassifications during counting and producing confidence intervals automatically from incoming data. To achieve this we avoid making best-guess pixel label decisions and opt instead for a probabilistically weighted count of evidence. This probabilistic approach also allows estimates to be written in a continuous form from which theoretical confidence intervals can be computed analytically, avoiding the need for Monte-Carlos.

Our estimate of the quantity of pixels associated with a given class of texture is given by:

$$Q(k) = \sum_d P(k|X_d) = \sum_d \frac{P(X_d|k)Q(k)}{\sum_j P(X_d|j)Q(j)}$$

where $Q(k)$ is the estimated number of pixels generated by class k; X_d is a data point taken from location d in an image; $P(k|X_d)$ is the probability that a data point was generated by class k; and $P(X_d|k)$ is the probability density for class k. A data point, X, describes a local image patch centred upon the pixel being classified. We have opted to use a BRIEF-like representation here using 18 random pixel pairs $[(\alpha_1, \beta_1), \ldots, (\alpha_{18}, \beta_{18})]$, within a 16 pixel radius. X takes the form of an 18 bit binary string with a 1 bit indicating pixel α_i is brighter than β_i and a 0 bit indicating otherwise. The density estimates, $P(X|k)$, are constructed from large histograms trained from example textures. An appropriate compact parametric form could be used instead for density estimation, but as a proof of concept we have opted for a simple histogram to avoid unnecessary complexity.

The maximum likelihood estimates for $Q(k)$ can be found by minimising the cost function:

$$F = -\sum_d \log \sum_k P(X_d|k)Q(k)$$

which can be achieved using the parameter update method of Expectation Maximisation.

Assuming density estimates match data being analysed and the EM algorithm converges successfully the estimates should be unbiased, i.e. the expectation of an estimate should be the true value. Statistical fluctuations in instantaneous estimates can be estimated by combining the independent variance contributions from the EM cost function and the Poisson errors found in density estimate histogram bins to give the covariance matrix:

$$\mathbf{C}_{ij} = cov_{CRB}(Q(i), Q(j)) + cov_{DE}(Q(i), Q(j))$$

where $cov_{CRB}(Q(i), Q(j))$ is the covariance of the cost function which can be found by inverting the Cramer Rao lower variance bound:

$$\mathbf{C}_{ij}^{-1} \approx \frac{\partial^2 F}{\partial Q(i) \partial Q(j)} = \frac{\sum_d P(i|X_d) P(j|X_d)}{Q(i) Q(j)}$$

The derivation of this lower bound assumes we are solving an Extended Maximum Likelihood problem, with the total quantity of data being Poisson distributed. Strictly, this should assume Binomial statistics, as the number of pixels in a image corresponds to a fixed number of events, N. The constraint on the number of pixels gives the more complex formula

$$\frac{\partial^2 F_{Binom}}{\partial Q(i) \partial Q(j)} = \sum_{d=1}^{N} \left(\frac{P(j|X_d)}{Q(j)} - \frac{1 - P(j|X_d) - P(k|X_d)}{N - Q(j) - Q(k)} \right) \left(\frac{P(k|X_d)}{Q(k)} - \frac{1 - P(j|X_d) - P(k|X_d)}{N - Q(j) - Q(k)} \right)$$

which we approximate with the much simplified Poisson version.

The additional contribution, $cov_{DE}(Q(i), Q(j))$, is the covariance due to density estimation errors, caused by the Poisson counting whilst populating histogram bins, computed using error propagation:

$$\begin{aligned} cov_{DE}(Q(i), Q(j)) = & \left[\sum_{k=i=j} \frac{Q(k)}{T(k)} \sum_{d \in R} P(\bar{k}|X_d)^2 \right] + \left[\sum_{k \neq i, k \neq j} \frac{Q(k)}{T(k)} \sum_{d \in R} P(i|X_d) P(j|X_d) \right] \\ & + \left[\sum_{k=i, k \neq j} \frac{Q(k)}{T(k)} \sum_{d \in R} -P(\bar{k}|X_d) P(j|X_d) \right] + \left[\sum_{k \neq i, k=j} \frac{Q(k)}{T(k)} \sum_{d \in R} -P(i|X_d) P(\bar{k}|X_d) \right] \end{aligned}$$

where $T(k)$ is the total quantity of training data used in the construction of the density estimate for class k; and $\bar{k}$ indicates all classes other than k.

The estimates of quantities and their variances can only be trusted if the modelled data densities match actual data being analysed. Any model-data mismatch will lead to systematic errors which grow as the amount of data analysed increases. Given the highly variable nature of Martian terrains, better density estimates can be produced by identifying and modelling multiple sub-textures, $k \in K$, to create multicomponent mixture models. The quantities of sub-textures can be summed to give overall estimates for each superordinate texture classification:

$$Q(K) = \sum_k Q(k)$$

When using mixture models the variances combine through error propagation as follows:

$$[var(Q(K))] = \mathbf{D}\mathbf{C}\mathbf{D}^T$$

where $\mathbf{C}$ is the covariance matrix defined previously, incorporating every sub-texture for each classification of Martian terrain; and $\mathbf{D}$ is a vector defining which sub-textures are relevant to classification K:

$$\mathbf{D} = [\delta_{1 \in K}, \delta_{2 \in K}, \ldots, \delta_{N \in K}]$$

with $\delta_{i \in K}$ equal 1 if subcomponent i belongs to class K, 0 otherwise.

The selection of sub-textures can be performed naively by simply subdividing a texture into a number of equally sized rectangles. Alternatively, sub-textures can be created systematically by clustering regions with similar density estimates. We have opted for a Bhattacharyya score as an appropriate measure of density estimate similarity [29]:

$$BC = \sum_X \sqrt{P(X|i) P(X|j)}$$

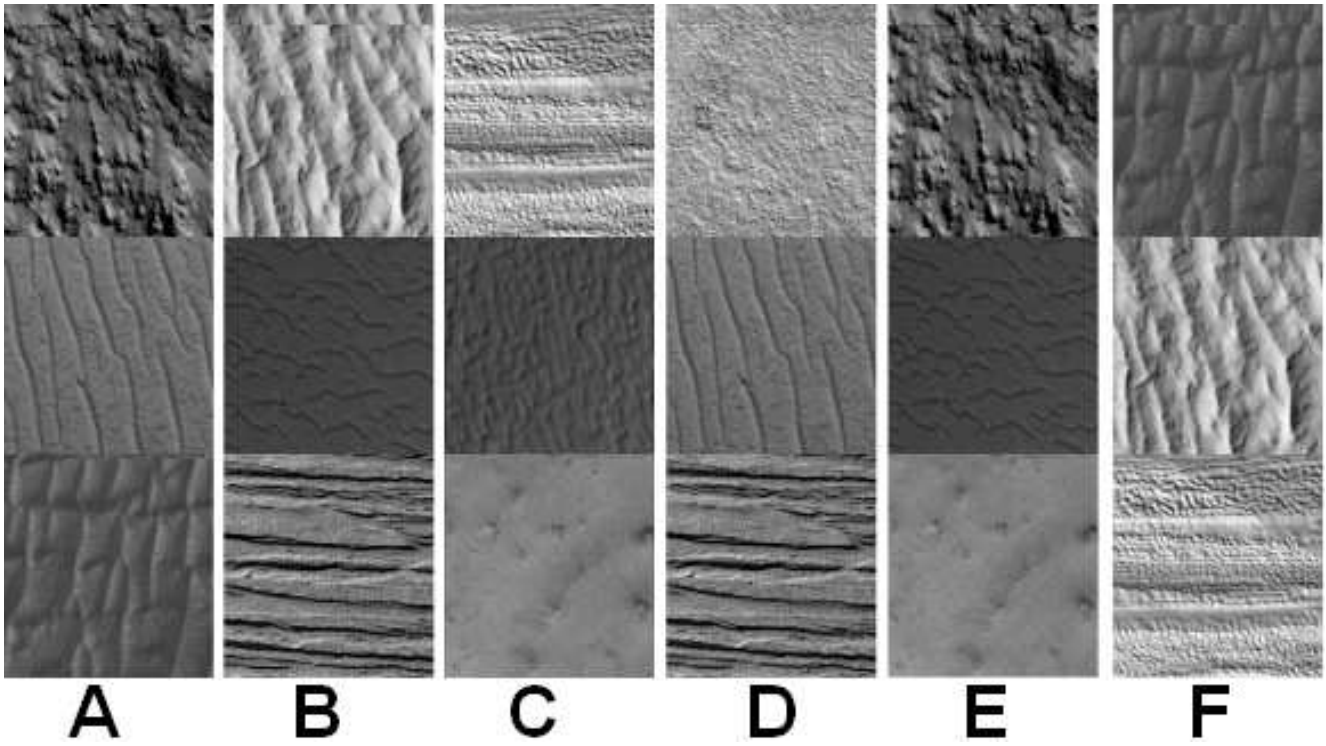

Figure 1: Samples of Martian textures taken from the datasets used in our experiments.

where $P(X|i)$ and $P(X|j)$ are density estimates generated from different regions of the same terrain texture. If the subregions are sufficiently similar they can be merged into a single density, otherwise they can form 2 distinctive sub-textures which can be linearly combined as part of a superordinate texture classification.

3 Experiments

The emphasis of our work is to produce quantitatively valid results for Martian images which fall within predictable confidence intervals. Our experiments reflect this by: a) focusing on actual Martian terrains rather than standard datasets; and b) by assessing results against theoretically computed variances, rather than against the performance of other algorithms.

10 Martian terrains were selected from a variety of MOC and HiRise satellite images from which 900x760 pixel regions were selected for training. Corresponding 300x760 pixel images were tiled together in 6 random groups of 3 forming test datasets A-F shown in Figure 1.

A series of experiments assessed the quantitative validity of our method by randomly analysing between 6250 and 200,000 pixels in each dataset. Each subtest of each experiment was performed 50 times, during which both theoretical and empirical errors were recorded. Theoretical errors were predicted using covariance matrices computed from incoming data. Empirical errors were determined by comparison to known ground truths. The experiments were divided into the following categories:

Single component with seen data

During these experiments density estimates were trained using a single component, i.e. only 1 sub-texture existed for each classification of terrain. The method was trained and tested using exactly the same data to ensure modelled densities perfectly matched the data densities observed during analysis. Under these circumstances we predicted that theoretically computed variances would exactly match those empirically observed.

Single component with unseen data

These experiments duplicated those previously undertake in 'single component with seen data' but with training and testing based on different examples of the same textures. This was intended to cause a mismatch between modelled and observed data densities. Under these circumstances we predicted that observed errors would become increasingly biased as the quantities of pixels analysed increased. The experiments were to reveal the severity of any bias.

Single component with partially seen data

These experiments tested the consistency of a texture by training and testing on the same data, but only testing on one half of the previously seen data. Although the data analysed had previously been seen during training, it was possible that the data density of the part did not match that of the whole. Under these circumstances we predicted that a single component model may not be sufficient for modelling textures causing a bias.

Mixture model with partially seen data

These experiments duplicated those previously undertaken in 'Single component with partially seen data', but with 2 sub-texture components naively trained using the left and right halves of the training image. Under these circumstances we predicted that when only half of the seen data was analysed the mixture model would ensure modelled and observed data densities matched, removing any bias, making theoretically computed and observed variances coincide just as they should under 'Single component with seen data'

Clustered mixture model with unseen data

These experiments used a series of mixture models of increasing complexity and applied them to unseen data. The training textures were divided into a coarse grid containing up to 20 cells per texture. Density estimates were constructed for each cell and densities which were sufficiently similar were merged to create sub-textures. The similarity threshold, based on the Bhattacharyya score, was gradually increased to enforce greater self-consistency within each sub-texture. In these experiments we predicted that as the similarity threshold became more stringent the number of necessary sub-texture components would increase, which would allows more accurate models to be created leading to a decrease in bias.

4 Results

Theoretical and empirical errors were compared by dividing residuals and standard deviations of actual errors by the standard deviations predicted by the covariance matrices. This produced a comparative mean and standard deviation for each subtest on each dataset and for each quantity of data analysed. If the errors were predicted correctly the comparative statistical error, i.e. standard deviation, should have an expectation of 1:

$$< \frac{\sigma_{observation}}{\sigma_{prediction}} >= 1$$

If there is no systematic error the comparative bias should have an expectation of 0:

$$< \frac{Q_{estimation} - Q_{truth}}{\sigma_{prediction}} >= 0$$

Figures 2 to 6 plot the statistical and systematic components of the errors for each category of experiment. Note that each individual data point represents the result of 50 iterations of each experiment and that results from all datasets are included in each plot.

5 Discussion

The method we have developed for quantifying the content of Martian images was driven from the outset by scientific requirements. Whilst absolute levels of performance in outputs are important, it is equally important

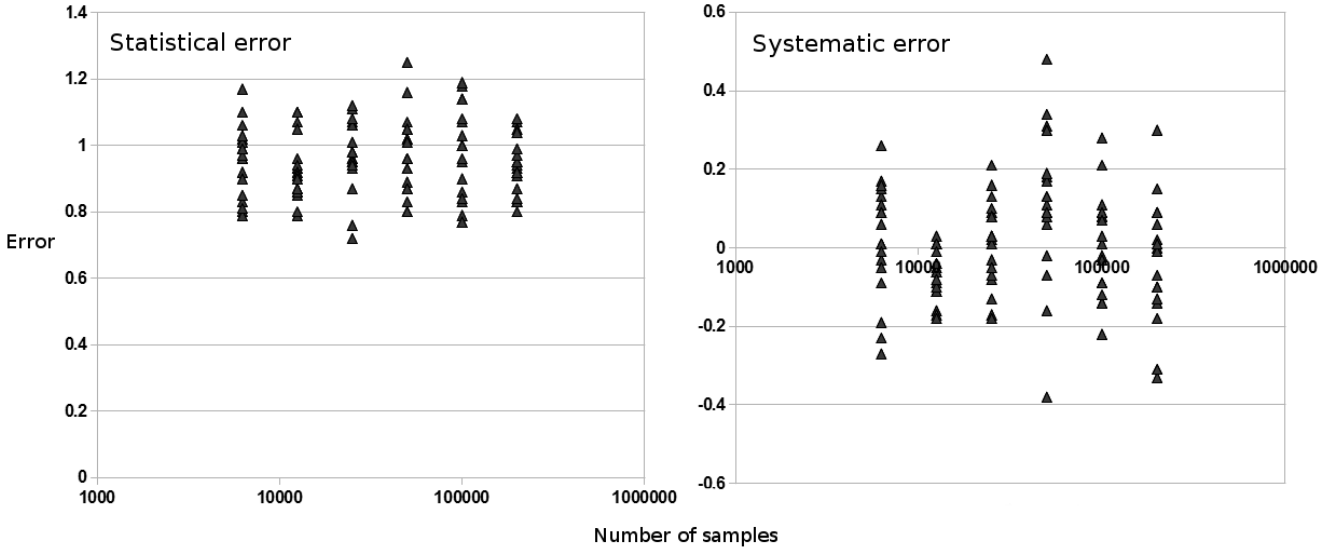

Figure 2: Single component with seen data. Here modelled data densities were forced to match actual data densities. Left: on average the spread of the observed error was correctly predicted. Right: on average there was no bias.

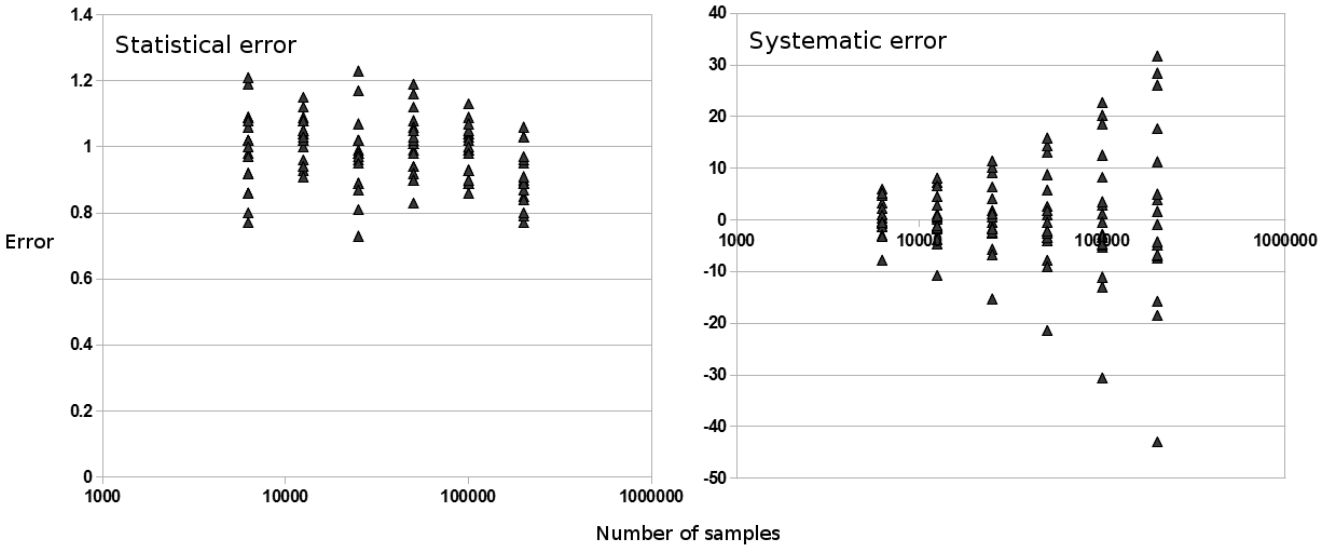

Figure 3: Single component with unseen data. Here modelled data densities did not perfectly match actual data as testing was performed on unseen images. Left: on average the spread of the observed error was correctly predicted. Right: as the quantity of data analysed increased a significant bias became evident.

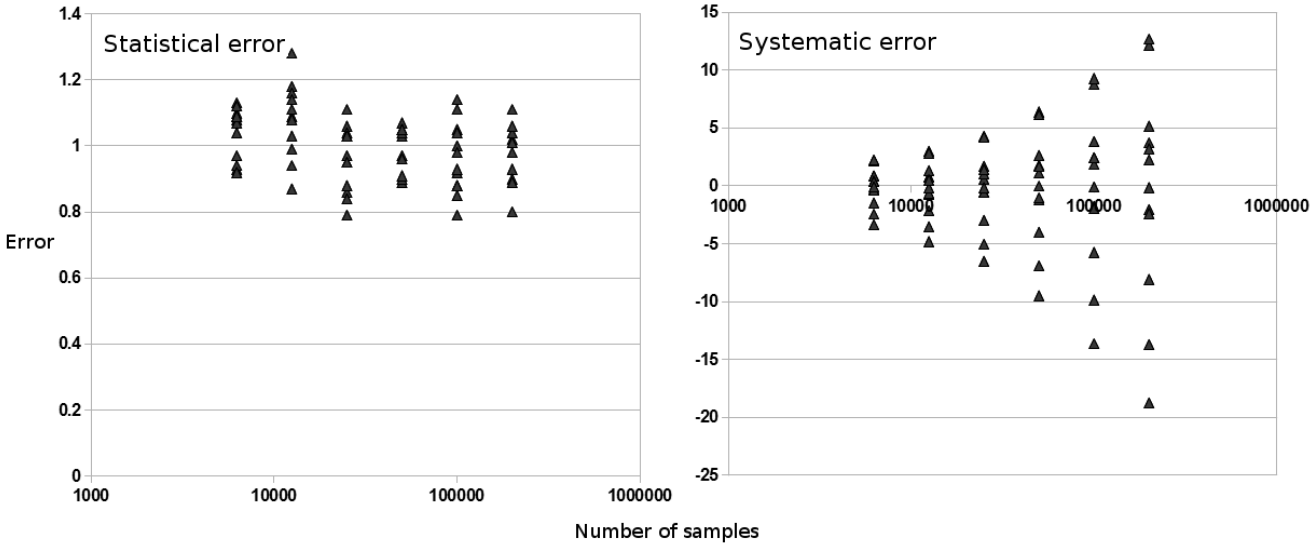

Figure 4: Single component with partially seen data. Here modelled data densities were applied to only part of the seen data. Left: on average the spread of the observed error was correctly predicted. Right: as the quantity of data analysed increased a bias was introduced, but not at the same rate as with completely unseen data.

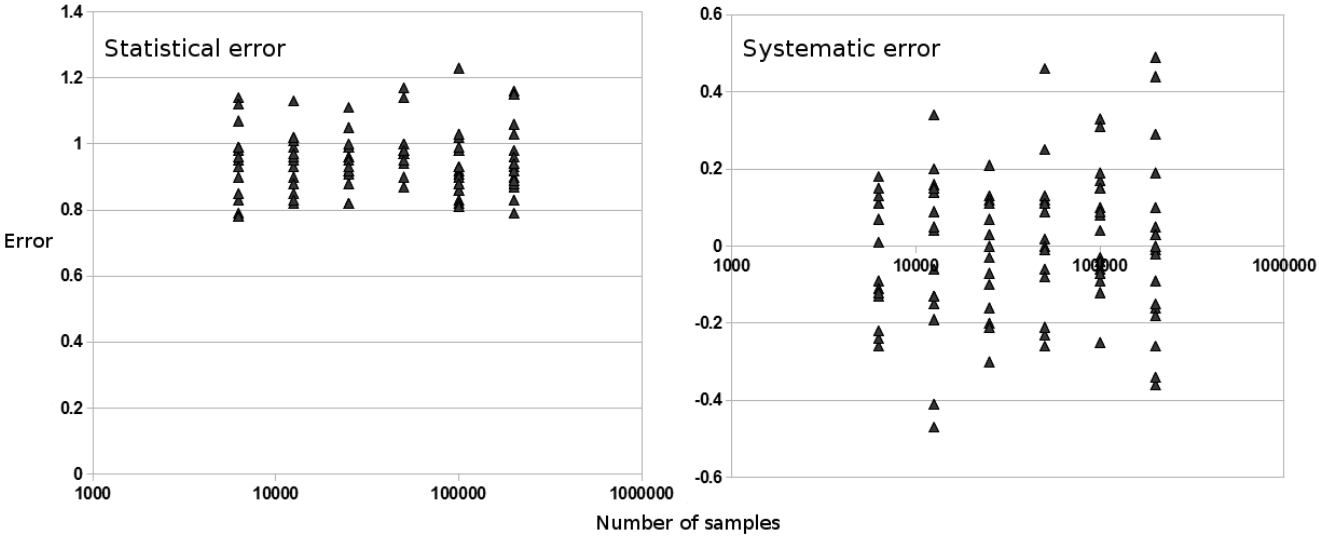

Figure 5: Mixture model with partially seen data. Here, a 2 component mixture model matched actual data densities. Left: on average the spread of the observed error was correctly predicted. Right: on average there was no bias.

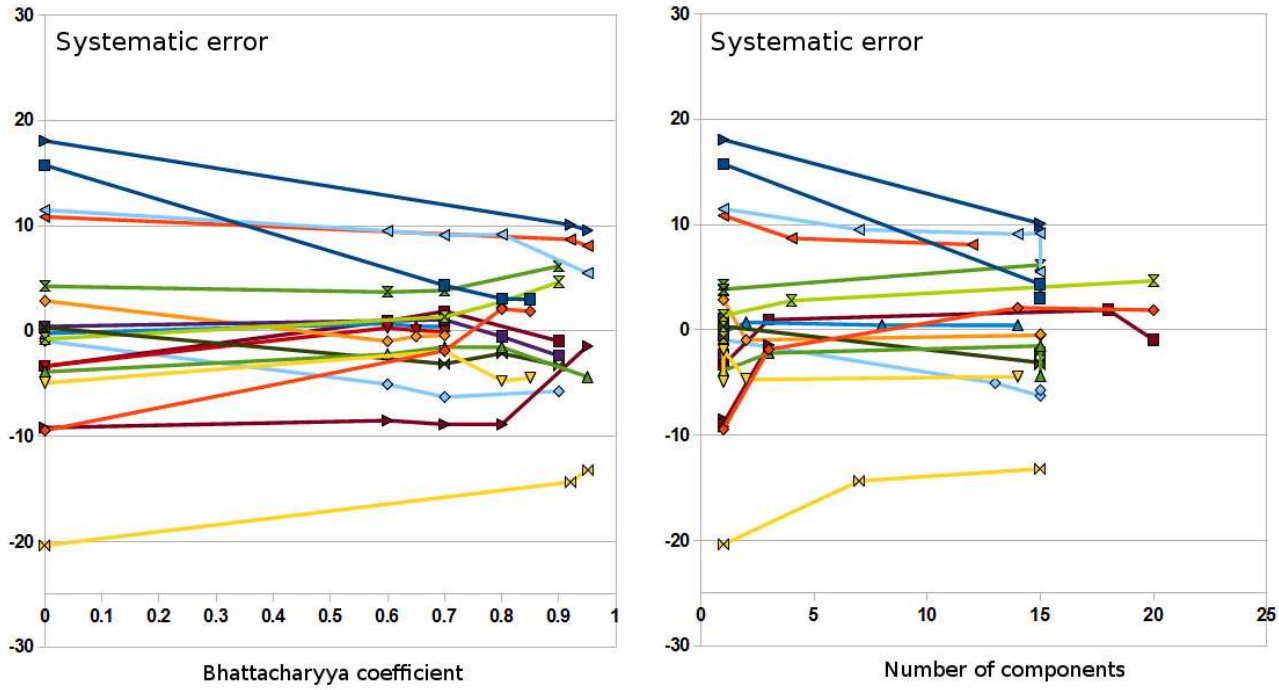

Figure 6: Clustered mixture model with unseen data. Here biases are plotted against different complexities of mixture models with the quantity of data analysed held fixed at 50,000 samples. Left: the similarity threshold, using the Bhattacharyya score, increases along the x-axis forcing sub-textures to be more self-similar. Right: the same data points are plotted but with the number of resulting sub-texture components along the x-axis. It can be seen that the more self-similar components are forced to be the more components are required to model them. It can also be seen that the bias can be reduced by up to a half using this simple clustering method.

for science to understand the certainty which can be placed in such outputs. A full covariance matrix provides a scientific researcher with the information required to understand the accuracy of results so they can be meaningfully compared against other results or theoretical models. Standard methods of classification such as SVM, RF, and Boosting etc. attempt to maximise absolute levels of performance but do not integrate this with any method to systematically provide quantitative confidence intervals. Our work demonstrates that it is possible to construct systems based upon well understood statistical and probabilistic theories (such as Bayes Theorem, the Cramer Rao Bound and error propagation) to provide confidence in outputs without the need of calibration with Monte-Carlo simulations. However, our methods will only work correctly if the density estimates generated from training data are truly representative of real unseen data analysed in practice. Our results show that for Martian terrains what a human may subjectively consider as being part of the same texture may objectively be significantly different causing systematic biases of up to 40 standard deviations when 200,000 pixels are analysed. This problem can be reduced somewhat by systematically determining appropriate subclasses of texture which can form mixture models better capable of matching real data densities, but these too have their limits in real data. With these issues in mind this discussion will try to answer the following questions: What is the best performance a general Bayesian classifier can achieve? How can this best performance be achieved? How does this best case relate to other classifiers? And in practice, how much confidence can be placed in our method?

Assuming quantities are converged upon using Expectation Maximisation, the best performance of a Bayesian classifier in estimating the total quantity of a given class in a dataset can be revealed by looking at the diagonal elements (i.e. variance) of the covariance matrix:

$$var_{CRB}(Q(k)) \approx \frac{Q(k)^2}{\sum_d P(k|X_d)^2}$$

If density estimations are perfect then there are no error contributions from $var_{DE}(Q(k))$, therefore in the best case only the above CRB contribution needs to be considered. The sum over the squared conditional probabilities will either increase or decrease depending on how probable a given class label is. The highest possible value of this sum is equal to $Q(k)$, and this is achieved when data points are completely unambiguous, i.e. for all X_d, the probability of k is either 0 or 1. In this case the sum amounts to counting up all the data points which are definitely from class k. The best case performance is therefore dominated by Poisson errors:

$$var_{best}(Q(k)) \approx \frac{Q(k)^2}{Q(k)} = Q(k)$$

which assumes the data generation process creating the k data points is a Poisson process with an expectation of $Q(k)$. The number of randomly selected ground truth pixels in our experiments were selected using a Poisson random number generator, but in real Martian images such Poisson processes exist naturally such as the quantity of impact craters expected to be found on a surface of a particular age. This Poisson best case accuracy also applies to other classifiers, such as Support Vector Machines and Random Forests, when all data points fall unambiguously on one side of a decision hyperplane or classification branch, respectively. However, when such clear-cut distinctions cannot be made there is no simple formula to express the uncertainty of quantities estimated using such classifiers, limiting their use in scientific applications, whereas our method allows covariances to be computed which account for additional uncertainty when data is ambiguous.

The best case performance can only be approached if the representation, X, is sufficiently detailed as to make density estimates, $P(X|k)$, highly distinctive from all other densities. Arguments of representational completeness can be used here to justify the selection of an encoding scheme. A complete representation allows the reconstruction of original data from the information coded within X. The Binary Robust Independent Elementary Features (BRIEF) style representation [31] we have employed has been demonstrated to outperform other encodings including SIFT [34] and SURF [35]. The original BRIEF work focused on matching features between images utilising large numbers of pixel pairs within extended regions and the Hamming distance as a similarity measure. Here, probability densities were built using thousands of small local image patches with relatively small numbers of pixel pairs. Despite the differences, the underlying representation of pairwise binary comparisons remains the same. Work presented by Kisilev et al. [33] shows how sets of pairwise comparisons can be used to reconstruct, up to a rank order, the shape of an underlying function. The same method could be applied to a binary representation to reconstruct image patches up to a rank order. This observation provides an explanation for the efficiency of an encoding which utilises a subset of random pixels to model a wider structure of stable relative pixel intensities. It must be noted that the quantitative validity of our method does not rely upon any one representation, however, the more discriminating the representation the lower the variance will be on estimated quantities. For science the ability to compute variances must override the desire for pure efficiency, but if efficiency is also required further research into discriminating representations and parametric density estimations may be necessary.

Our method will only operate correctly if the training provided is truly representative of real data. Experiments show how severe bias can become in cases where there is a mismatch between modelled data densities and actual data, raising the question of how useful the method is for real science. It is acceptable for the bias to be up to half of the statistical error, in which case the predictable component is still the dominant source of uncertainty. Beyond this point the method can not be trusted. The bias increases with the quantity of data analysed, so it is possible to limit usage to smaller datasets. Only applying the method at edges rather than at each pixel will instantly increase the overall size of applicable images. Training can also be improved by providing more examples and constructing more sophisticated mixture models which cluster sub-textures into all possible views of the data. Experiments have shown that the mixture modelling approach will reduce bias. Extrapolating these results, in theory it should be possible to construct models which bring the bias below half of the statistical error, given the right training.

6 Conclusion

We have demonstrated the possibility of creating a quantitative system for the scientific analysis of Martian terrains with a strong emphasis on providing not only raw outputs but also covariances which allow outputs to be interpreted with confidence. This is in contrast to alternative systems which focus largely on maximising absolute performance but do not attempt to provide any indication of how much outputs can be trusted. Our system will work reliably as long as training data matches data analysed in practice. If there is disagreement between modelled data densities and actual data then systematic effects can grow to dominate results, but methods are available to reduce biases or at least warn a user of their presence. Whilst our focus is the analysis of Martian images our method could be applied to other texture analysis tasks. More work is required to generalise our method to allow it to be used on multiple scales and orientations of Martian terrains.

7 References

1. McEwen AS et al. (2007), Mars Reconnaissance Orbiter's High Resolution Imaging Science Experiment (HiRISE), Journal of Geophysical Research, vol 112, E5
2. Hawkins SE et al. (2007), The Mercury dual imaging system on the MESSENGER spacecraft, Space Science Reviews, vol 131, 1-4, 247-338
3. Russell C et al. (2007), Dawn Mission to Vesta and Ceres - Symbiosis Between Terrestrial Observations and Robotic Exploration, Earth Moon and Planets, vol 101, 65-91
4. Blum M D et al. (2000), Fluvial responses to climate and sea-level change, Sedimentology, vol 47, 2-48
5. Head J W et al. (2003), Recent ice ages on Mars, Nature, vol 426, 6968, 797-802
6. Wei Luo et al. (2009), Journal of Geophysical Research, vol 114, E11010
7. Silvestro S et al. (2011), Present-day Aeolian Activity in Arabia Terra, first results from a global mapping of active dune fields on Mars, 42nd Lunar and Planetary Science Conference
8. Neukum G (2001), Cratering records in the inner solar system in relation to the lunar reference system, Space Science Reviews, vol 96, 55
9. Plesko S C et al. (2006), A Statistical Analysis of Automated Crater Counts in MOC and HRSC Data, 37th Lunar and Planetary Science Conference
10. Svein O K et al. (2007), Automatic and semi-automatic detection of possible meteorite impact structures in the Fennoscandian shield using pattern recognition of spatial data, Proceedings ScanGIS, 227-235
11. Smirnov A (2002), Exploratory Study of Automated Crater Detection Algorithm, Dept. of Computer Science, University of Colorado, Internal Report, www.cs.colorado.edu / rossnd /fcdmf /CraterPaper.pdf
12. Stepinski T (2006), Machine Detection of Martian Craters from Digital Topography, 37th Lunar and Planetary Science Conference, LPI
13. Salamuniccar G et al. (2006), Estimation of False Detections for Evaluation of Crater Detection Algorithms, 37th Lunar and Planetary Science Conference
14. Kim J et al. (2004), Quantitative Assessment of Automated Crater Detection on Mars, 20th ISPRS Congress
15. Magee M et al. (2003), Automated Identification of Martian Craters Using Image Processing, 34th Lunar and Planetary Science Conference
16. Earl J et al. (2005), Automatic Recognition of Crater-like Structures in Terrestrial and Planetary Images, 36th Lunar and Planetary Science Conference
17. Honda R et al. (2000), Data Mining System for Planetary Images: Crater Detection and Categorization, Proceedings of the International Workshop on Machine Learning of Spatial Knowledge, ICML, 103-108
18. Martins R et al. (2008), Crater Detection by a Boosting Approach, Geoscience and Remote Sensing Letters, vol 6 i 1, 127-131
19. Ding w et al. (2011), Sub-Kilometer Crater Discovery with Boosting and Transfer Learning, ACM Transactions on Computational Logic
20. Tar P (2010), The Application of Appearance Models to Martian Impact Craters, The University of Manchester, Internal Report:Tina memo 2010-011, www.tina-vision.net /docs/memos/2010-011.pdf
21. Stepinski T (2004), Extraction of Martian Valley Networks from Digital Topography, Journal of Geophysical Research Vol 109, page 9
22. Stepinski T (2006), Properties of Martian Highlands Drainage From Themis Image and MOLA Topography, 37th Lunar and Planetary Science Conference
23. Vaz DA et al. (2006), Automatic Detection and Classification of Fault Scarps on MOLA data, 37th Lunar and Planetary Science Conference
24. Gaetano R et al. (2010), Graph-based Analysis of Textured Images for Hierarchical Segmentation, BMVC 2010
25. Jia W et al. (2010), Classifying Textile Designs using Region Graphs, BMVC 2010
26. Heess N et al. (2009), Learning Generative Texture Models with Extended Fields-of-Experts, BMVC 2009
27. Santner J et al. (2009), Interactive Texture Segmentation using Random Forests and Total Variation, BMVC 2009
28. Corani G et al. (2010), Robust Texture Recognition Using Credal Classifiers, BMVC 2010
29. Thacker N.A. et al. (1997), The Bhattacharyya Metric as an Absolute Similarity Measure for Frequency Coded Data, Kybernetika, 34, 4, 363-368
30. Viola P, Micheal J (2001), Robust Real-time Object Detection, Second International Workshop on Statistical and Computational Theories of Vision
31. Calonder M et al. (2010), BRIEF: Binary Robust Independent Elementary Features, EPFL, ECCV conference paper

32. Brodatz P (1966), Textures: A Photographic Album for Artists and Designers, Dover Publications, ISBN 0486406997
33. Kisilev, Pavel and Freedman (2010), Parameter Tuning by Pairwise Preferences, Proceedings of the British Machine Vision Conference, BMVA Press
34. Lowe D (2004), Distinctive Image Features from Scale-Invariant Keypoints, Computer Vision and Image Understanding 20, 91-110
35. Bay H et al. (2008), SURF: Speeded Up Robust Features, Computer Vision and Image Understanding 10, 346-359