TEXTURECAM: AUTONOMOUS IMAGE ANALYSIS FOR ASTROBIOLOGY SURVEY David R. Thompson[1,2], Abigail Allwood[2], Dmitriy Bekker[2], Nathalie A. Cabrol[3], Tara Estlin[2], Thomas Fuchs[4], Kiri L. Wagstaff[2]; [1]david.r.thompson@jpl.nasa.gov; [2]Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Dr. Pasadena, CA 91109, USA; [3]SETI Institute, Mountain View, CA 94043, USA; [4]California Institute of Technology, Pasadena CA 91125

Introduction: The TextureCam concept is a "smart camera" that aims to improve scientific return by increasing science autonomy and observation capabilities both when the spacecraft is stationary and when it is traversing. This basic onboard science understanding can summarize terrain encountered during travel [1] and direct autonomous instrument deployment to targets of opportunity [2]. Such technology would assist Mars sample caching or return missions where the spacecraft must survey vast areas to identify potential sampling sites and habitat indicators.

The instrument will use texture channels to differentiate and map habitat-relevant surfaces. Here we use the term *texture* not in the geologic sense of formal physical properties, but rather in the computer vision sense to mean statistical patterns of image pixels. These numerical signatures can in turn distinguish geologically relevant elements such as roughness, pavement coatings, regolith characteristics, sedimentary fabrics and differential outcrop weathering. Similar algorithms can perform microstructure analysis and sedimentology using microscopic images. On the scale of meters, surfaces and features can be recognized to summarize rover traverse and draft surficial maps for downlink.

Reliable image classification in field conditions is difficult due to factors like variable lighting, surface conditions, and sediment deposition. Future development aims for reliable recognition of basic geological elements. Rad-hardened FPGA processing will instantiate these general-purpose texture analysis algorithms as part of an integrated, flight-relevant prototype.

Experimental setup: An initial test demonstrates automatic recognition of stromatoform structures in outcrop. These surfaces have distinctive laminar textures that may indicate previous biogenic activity [3] and would be valuable targets for followup investigation and human review. While it would be unlikely to find such features exposed on the Mars surface, they provide a controlled test to guide initial development and evaluation. We use a machine learning approach, training the system on example pixel labels supplied by the designer. This yields a statistical model of pixel relationships that can classify new surfaces (Figure 1).

Stromatolite images were acquired from the Pilbara formation [3]. We consider three outcrops of laminar structures with *conical* and *wavy* morphologies. We also include a scene containing much finer-scale localized *egg carton* laminations. These all show a common lay-
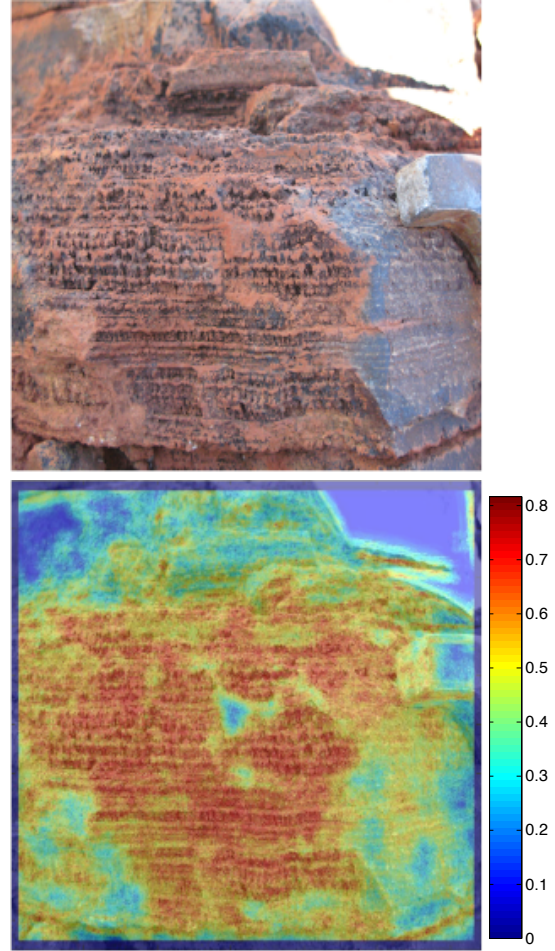

Figure 1: Top: Original image. Bottom: Automatic surface map showing the probability of belonging to the *stromatolite* class. Green, yellow and orange are uncertain classifications.

ering structure that we aim to distinguish automatically from the surrounding surface. We hypothesize that the statistical properties of the image textures are generalizable across the three stromatolite types.

Method: All image pixels were labeled by hand as belonging to one of three classes: (1) stromatoform, (2) non-stromatoform, or (3) ambiguous/unclassified. A high-pass filter mitigates most differences in shadow and illumination and reveals small-scale surface texture (Figure 2 Right). We use a single color channel, and rotate the training image and labels by 15-degree increments to provide generality and rotational invariance.

We use this training data to construct a *random forest* pixel classification system [4]. It combines the result of

Figure 2: Left: Training image, a conical stromatolite. Right: A high-pass filter reduces variability due to shadows and illumination.


Figure 3: Decision tree classification. We aim to classify the pixel marked with an 'X;' it propagates down a tree of left/right decisions to arrive at a classification label. The decision at each intermediate node applies a threshold test to one or more pixel values from its local neighborhood.

multiple independent *decision tree* classifiers. Each tree is a hierarchical sequence of tests applied to local image values in the neighborhood of the classified pixel. Each test considers a numerical attribute such as (1) the absolute intensity of a nearby location, at some offset relative to the target pixel; (2) the difference in intensity between two nearby locations; or (3) the absolute difference in intensity between two nearby locations. A threshold on the attribute determines whether the pixel will propagate left or right. Such tree sequences can form complex decisions based on local edges, discontinuities, and other characteristic microtexture. Figure 3 illustrates the concept for classification of a pixel marked X. It propagates down a simple tree of binary decisions - tests on values within a local neighborhood represented by the red rectangle. The population of pixels arriving at each final node of the decision tree gives a distribution over classes for pixels that follow this path.

The training procedure grows each decision tree from a single root node. At each new decision node, it searches a random subset of possible tests to find the most informative one as determined by an *expected information gain* metric [5]. This score captures the un-
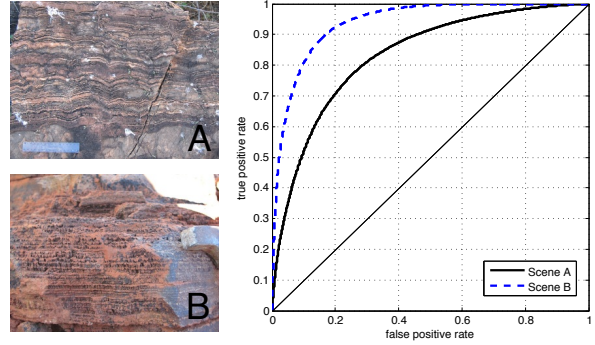

Figure 4: Test scene ROC performance

certainty in classes of the pixels in the resulting left and right subpopulations, weighted by size. We grow all trees for a fixed number of iterations. The process takes just a few minutes on a modern consumer-grade CPU.

Results: We evaluate performance on each image by training on the two other images (Leave One Out Cross Validation). Figure 4 shows Receiver Operating Characteristic (ROC) curves for the pixels in two scenes: scene A with wavy morphology; and B which introduces weathering and variable depth-of-field. The image of a conical stromatolite (Figure 2) has insufficient stromatoform pixels for a confident ROC estimate so we omit it from the test set. ROC performance suggests that the laminar structures are recognizable across images, with reasonable performance given the limited training data.

Figure 1 shows the original image for scene B along with the final surficial classification probabilities. Values near $P = 0.5$ indicate ambiguous classifications. The classification learns to favor areas with multiple narrow layers; aligned fractures in the non-stromatolite surface are a main source of errors. One solution could be a larger training set that includes non-stromatolite laminae. A complementary solution is a subsequent segmentation step to clean the map into contiguous areas and eliminate small clusters of outlier pixels without support from neighboring areas [5].

References: [1] D. Thompson, et al. (2011) *Journal of Field Robotics* 28(4):542. [2] T. Estlin, et al. (to appear) *ACM Trans on Intelligent Systems and Technology*. [3] A. Allwood, et al. (2006) *Nature* 441(7094):714. [4] L. Breiman (2001) *Machine learning* 45(1):5. [5] J. Shotton, et al. (2008) *IEEE Conf on Computer Vision and Pattern Recognition* 1–8.