TextureCam: A Smart Camera for Microscale, Mesoscale, and Deep Space Applications. William Abbey[2], Abigail Allwood[2], Dmitriy Bekker[2], Benjamin Bornstein[2], Nathalie A. Cabrol[3], Rebecca Castaño[2], Steve A Chien[2], Joshua Doubleday[2], Tara Estlin[2], Greydon Foil[4], Thomas Fuchs[2], Daniel Howarth[4], Kevin Ortega[2], David R. Thompson[1,2], Kiri L. Wagstaff[2]. [1]Contact: 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] Carnegie Mellon University. 5000 Forbes Ave., Pittsburgh, PA 15213, USA.

Introduction: The TextureCam project is developing a "smart camera" that can classify geologic surfaces in planetary images. This would allow autonomous spacecraft to collect data opportunistically during intervals between communications with Earth, such as during long traverses [1,2]. Its surface classifications can identify new targets that were not anticipated in advance. The spacecraft might use this information to target these features with high-resolution instruments such as spectrometers and narrow-field cameras. Classifications could also inform data "triage" decisions, identifying high value images for prioritized downlink. Finally, the surface classifications can serve as compressed maps of image content. Each of these strategies can improve the science data returned at each command cycle and speed reconnaissance during site survey and astrobiology investigations. Our first year of development has completed the image analysis algorithms and validated them in software tests. Here we survey these initial results, and explore several application areas relevant to Mars and beyond.

Method: The TextureCam instrument consists of a Virtex 5 FPGA component connected to any framing camera. Our prototype uses a commercial camera with a high-speed CameraLink interface. The image analysis is currently implemented in software only. We use a *random forest* classification engine similar to Shotton et al [4], and fit this model using labeled examples supplied in advance by the analyst. The training process optimizes a "decision tree," a sequence of simple numerical tests applied to local neighborhoods of the image that together determine the texture classification of each pixel. This procedure is detailed in Thompson et al. [5]. After training, the model can extrapolate the statistical relationships to classify other scenes. Table 1 shows the expected specifications of our prototype.

Microscale Geologic Surface Classification: We have tested the texture retrieval software on two microscale imaging applications. First, we demonstrated simple particle size classifications into Wentworth size categories (Figure 1). These would permit rovers to interpret the sedimentology of the terrain en route, and could assist analyses by scientists on the ground. The image shows a simple size classification of a Mars Exploration Rover Microscopic Image into very coarse and medium sand categories (blue, green). We have also tested the system in the laboratory, recovering image textures of smooth-cut rock samples imaged under diverse lighting conditions. Tests of basalt samples suggest good performance as long as the lighting direction remains within approximately 45 degrees of vertical (i.e. midday imaging conditions – Figure 5).

Mesoscale Geologic Surface Classification: Recent tests on images collected in the Mojave Desert demonstrate texture classification performance that is invariant to range. We separate the image into near-field, mid-range, and far-field segments (0-2m, 2-6m, and >6m respectively), training a separate classifier model on each subset. At runtime, stereo data identifies the parts of the scene which are relevant for each classifier. Cast shadows are apparent at mesoscales; we ignore these areas with a simple intensity threshold. Figure 5 shows the result on a typical scene from the Mojave Desert.

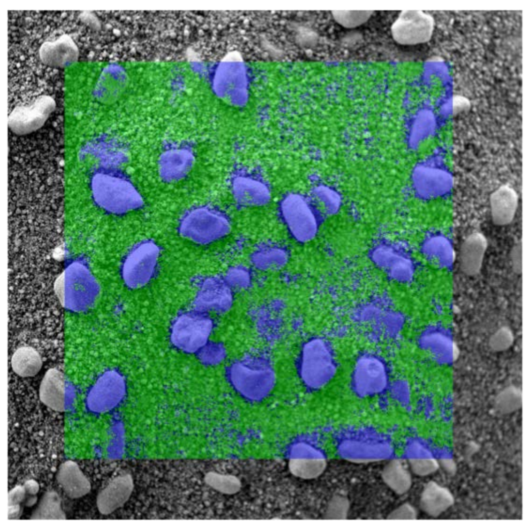

Figure 1. Automatic inference of particle sizes in MI images. Here the system identifies parts of an image whose particle sizes are predominantly large sand (blue) and medium sand (green).

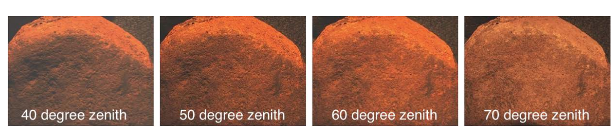

Figure 2. Laboratory imaging demonstrates surface texture variability as a function of lighting angle.

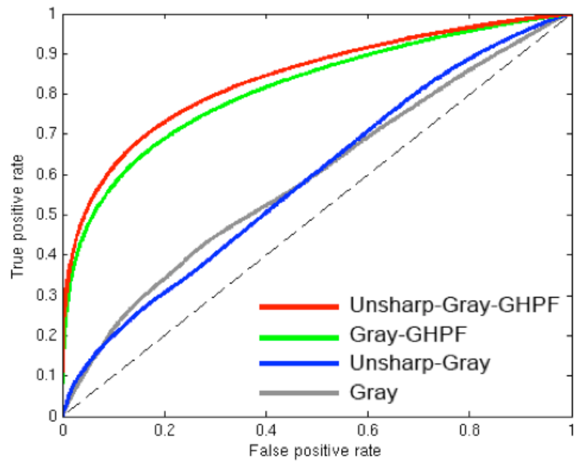

Figure 3. Preprocessing with a "high pass filter" (green and red lines) removes large-scale shading effects and improves model performance under variable illumination.

Deep Space and Orbital Applications: The same techniques can be used at larger scales. We are incorporating the methods into a flight demonstration on the IPEX cubesat, a microsatellite scheduled for launch in Oct. 2013 [6]. The spacecraft will carry a simple RGB framing camera, and periodically downlink images of the Earth. However, the bandwidth of communication with this spacecraft is limited so only one or two images can be downliked each day. The random forest classifier will run onboard the serial CPU, where its results can identify the best images for downlink. We will attempt to classify image regions corresponding to four classes: the planetary limb, deep space, clear terrain, and cloudy terrain. Downlink will favor images containing as much cloud-free terrain as possible. As a side benefit the model will generate compressable ~2KB "quicklook" maps of the scene content for downlink at very little bandwidth cost (Figure 5). A recent blind test on images from two high-altitude balloon flights demonstrate that our initial model generalizes well to new scenes. If successful, these tests will pave the way for further use of intelligent science image analysis onboard spacecraft.

Acknowledgements: A portion of this research was carried out at the Jet Propulsion Laboratory, California Institute of Technology. Copyright 2013 California Institute of Technology. All Rights Reserved; U.S. Government Support Acknowledged. The TextureCam project is supported by the NASA Astrobiology Science and Technology Instrument Development program (NNH10ZDA001N-ASTID).

References: [1] D. Thompson, et al. (2011) *Journal of Field Robotics* 28(4):542. [2] T. Estlin, et al. *ACM Trans on Intelligent Systems and Technology.* [3] L. Breiman (2001) *Machine learning* 45(1):5. [4] J. Shotton, et al. (2008) *IEEE Conf on Computer Vision and Pattern Recognition* 1–8. [5] D. R. Thompson, et al. (2012) *i-SAIRAS.*

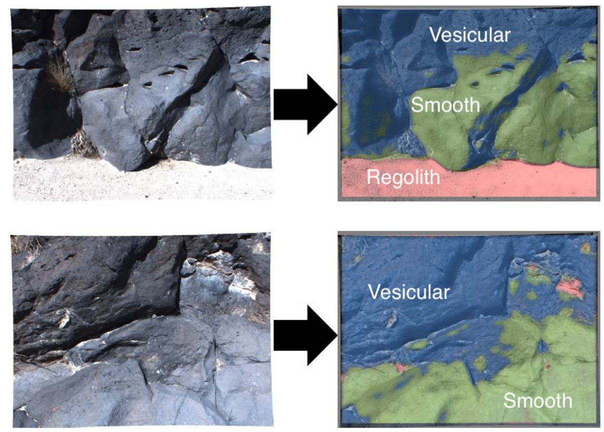

Figure 4. Mesoscale scene from the Cima Volcanic Fields in the Mojave Desert. It contains vesicular basalt (blue), smooth basalt (green), and unstructured regolith (red). The left side shows the initial image. We train the model on separate images and applying it to this scene, yielding automatic classifications at right.

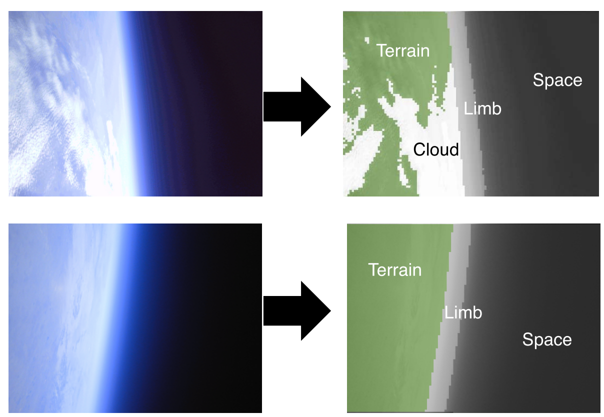

Figure 5: Images from high-altitude balloon flights demonstrate a simple orbital mapping task. Top: training image from first balloon flight. Bottom: test image from second balloon flight.

	Specification
Pixel resolution	1200x1600px
Classification frame rate	1Hz
Power consumption	~30W (2 cameras, FPGA)
Operating range	6cm-Infinity

Table 1: System specifications for year 2 field prototype