FEATURE DETECTION ONBOARD MARS ROVERS: AUTOMATED CLOUD AND DUST DEVIL DETECTION

A. Castano[1], A. Fukunaga[1], R. Castano[1], S. Chien[1], R. Greeley[2], P. Whelley[2], L. Neakrase[2], and M. Lemmon[3], [1]Jet Propulsion Laboratory, Pasadena, CA 91109, andres.castano@jpl.nasa.gov, [2]Arizona State University, Tempe, AZ, 85287, [3]Texas A&M University, College Station, TX, 77843.

Introduction: Recent explorations on the Martian surface have revealed an environment far more dynamic than previously believed. In particular, the atmosphere of Mars is very dynamic. Dust devils and clouds are dynamic atmospheric features that have been observed by the Mars Exploration Rovers (MER). These high science value events have been the subject of considerable study. Both dust devil and cloud detection campaigns have been conducted, but in general these are rare events. For example, only around 10-25% of the cloud campaign images collected have clouds in them. Prior campaigns have involved collecting images at fixed times for return to Earth. This is an inefficient use of downlink bandwidth as the majority of images do not contain dust devils or clouds.

To improve the effectiveness of atmospheric imaging campaigns, we have developed a different approach. In this approach onboard processing is used to screen images for the science features of interest (i.e., clouds and dust devils). Using this approach, many images can be collected onboard resulting in a much greater time range for capturing the rare phenomena. Even when the images cannot be down-linked (such as when too many events are detected), compact summary statistics on the number and type of events can still be down-linked to provide valuable information. The code has been integrated with the MER flight software, and is scheduled for upload to the MER rovers as part of the R9.2 software upgrade, pending final regression testing [1,2].

Dust Devil Detection: The two most common methods for detecting dust devils are the comparison of two or more spectral bands of the scene and the motion detection using a temporal sequence. We have selected the latter as it has application to Pancam, Navcam and Hazcam imagery. In theory, detecting motion in the scene is not equal to detecting dust devils. In practice, if image noise can be accounted for, the majority of changes in a sequence of images taken over a short time period of a scene on Mars will be from dust devils.

The challenge for robust automated detection occurs when the difference in the intensity of the two images, at the location of the change, is comparable in magnitude to the noise of the image. For such situations, a discriminating threshold cannot be selected easily as it will invariably consider image noise as change (false positive), actual change as noise (false negative) or both. To reduce the noise, we detect changes in image I_i using the average of *n* images of the sequence, I_o, and I_{io}, the average of the *n-1* images of the sequence that excludes I_i, i.e.,

$$I_o = \frac{1}{n}\left(\sum_{j=1}^{n} I_j\right) \text{ and } I_{io} = \frac{1}{n-1}\left(\sum_{j=1}^{n} I_j - I_i\right).$$

The difference of these images contains the average of the image noise for all the areas where I_i was equal to the other images and the average of the image noise plus the change for the areas where I_i was different to the other images of the sequence, i.e., the intensity of the change is damped by a factor of *1/n*. Assuming that the major component of the image noise is zero-mean Gaussian noise, then the areas with no change tend to zero while the areas with change do not. Thus, although the intensity of the motion information has been damped, the motion can be detected because the areas with no change tend to zero faster than those with change. To complement this approach we use a minimized version of the image, a threshold biased by the local noise, and blob filters.

The algorithm was tested on 25 image sequences, all acquired on Mars using the left Navcam of the Spirit rover. Each sequence had a length that varied between 6 and 20 images. The set of sequences was biased toward faint dust devils. Given these sequences, we analyzed all the possible subsets of a given number of contiguous images for 4, 6, and 8 consecutive images. The results are in Table 1. Figure 1 is a scene with several dust devils detected.

Cloud Detection: The approach used to detect clouds is to assume that large variations in the intensity of the sky in the image should correspond to clouds; this assumption holds true with the exception of large changes of local intensity due to zero-mean Gaussian noise (e.g., particularly noticeable at dusk and dawn) and large changes of global intensity due to camera effects (e.g., vignetting-like effect that darkens the corners of the image under low-light conditions). The first step to analyze the sky intensity is to segment the image using the sky detector previously developed under OASIS [3]. The result of the sky detector is used to mask out the ground and, if desired, to buffer an area above the skyline, which avoids illumination effects frequently found near the horizon.

Once the sky has been segmented, we search for changes in the sky by using an edge detector; strong edges indicate large gradients on the sky that are caused by the presence of clouds. The threshold that determines the value of the edge that corresponds to a cloud is weighted by the noise of the image.

The algorithm was tested using a set of 210 images taken on Mars by Spirit and Opportunity. All of the images contained the sky, and most of the images contained both sky and ground. 47 of the images were images that a MER scientist had labeled as containing clouds, while the remaining 163 images were selected randomly from the set of all MER images that contained the sky, and manually verified as not containing clouds. The images with clouds were further divided into 29 images that contained evident clouds, 13 images that contained soft, hard-to-see wispy clouds and 5 images for which the scientists could not decide if there was a cloud or not. For this set, the algorithm detected 100% of the evident clouds, 100% of the wispy clouds and 60% of the ones in the undecided subset. Likewise, it stated correctly that were no clouds in the no-cloud set 93.2% of the time. In summary, the algorithm was correct at identifying whether or not a cloud was present in 93.3% of the test set; there were 3 false negatives and 11 false positives. A sample image and the clouds that were detected are shown in Fig. 2.

References: [1] Castano, et al., *IEEE Aerospace*, (2006). [2] Schenker, et al, *IEEE Aerospace*, (2006). [3] Anderson, R. C., et al., *Fall AGU*, (2003).

Figure 1. Result of motion detection in an image. Two of the dust devils are evident (3rd and 5th red box), while the other three are difficult to see without the motion sequence.

Table 1. Results from dust devil detection algorithm.

No. images	n-tuples	+	-	correct	False -	False +
4	279	120	159	237 (84.9%)	10	32
6	228	121	107	190 (83.3%)	14	24
8	180	116	64	155 (86.1%)	10	15

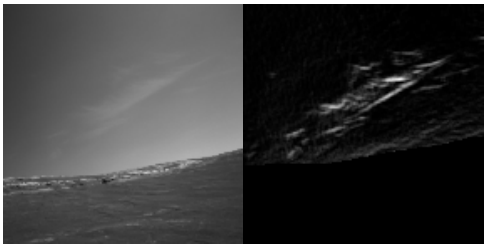

Figure 2. An example of cloud detection. Left image is the original image and the right image is the result of the cloud detection algorithm.