SCIENCE-BASED REGION-OF-INTEREST IMAGE COMPRESSION. K. L. Wagstaff, R. Castano, S. Dolinar, M. Klimesh, R. Mukai, *Jet Propulsion Laboratory, Pasadena CA 91109, USA, (kiri.wagstaff@jpl.nasa.gov)*.

As the number of currently active space missions increases, so does competition for Deep Space Network (DSN) resources. Even given unbounded DSN time, power and weight constraints onboard the spacecraft limit the maximum possible data transmission rate. These factors highlight a critical need for very effective data compression schemes. Images tend to be the most bandwidth-intensive data, so image compression methods are particularly valuable. In this paper, we describe a method for prioritizing regions in an image based on their scientific value. Using a wavelet compression method that can incorporate priority information, we ensure that the highest priority regions are transmitted with the highest fidelity.

There are three parameters that affect the level of effective compression: the raw data acquisition rate α, the internal buffer size β, and the data transmission rate τ. If λ is the lossless compression coefficient (typically 0.4 to 0.7), then when $\tau \geq \lambda\alpha$, all of the collected data can be transmitted without loss. Otherwise, some data must be discarded. The compression software being used by the Mars Exploration Rovers, ICER [4], performs wavelet-based compression and prioritizes bit layers based on their contribution to overall image quality. The successor to this technology is ROI-ICER, which prioritizes compressed data based on region-of-interest information [3]. ROI-ICER allocates more transmission bits to areas of the image designated as high-priority. How to best determine the relative priorities of different parts of an image is an open question, and generally must rely on domain-specific information. For example, Dolinar et al. [3] calculated priorities based on temperature information for Earth images to detect information about forest fires.

Science Goals. In this paper, we focus on methods for calculating priority information for images collected by Mars rovers. In particular, we seek to automatically prioritize regions of an image that contain rocks over those that do not. Image areas that contain sand, rover parts, etc. are useful for providing context, but we would prefer to sacrifice detail in those areas of the image in favor of preserving details of the rocks.

The images used in this study were taken by the FIDO (Field Integrated Design and Operations) rover, a field test rover for the Mars Exploration Rovers. The images are greyscale views of a field test area near Flagstaff, Arizona. Each image is 640 $\times$ 480 8-bit pixels and 300 kilobytes in size. The full set of 25 images contains a total of 185 rocks.

Creation of Priority Maps. Our goal is to prioritize image regions that contain rock information. For these initial experiments, we manually identified the rock boundaries. In the near future, we expect to incorporate the automated rock finder that has been developed at JPL [1]. This rock finder is successfully being used to analyze recent images returned by Spirit, the Mars Exploration Rover that landed on January 3, 2004 (PST).

ROI-ICER allows the specification of up to eight priority levels, but we need only two for our purposes. Therefore, we generated a priority map for each image with rock regions marked as priority 7 and non-rock regions marked as priority 0. This corresponds roughly to indicating that rock pixels should be represented with seven more bits of precision than non-rock pixels, thereby increasing the degree to which the rock pixel values are preserved during compression.

Experimental Setup and Evaluation. We conducted a series of experiments to determine the effectiveness of our science-based priority maps. We simulated the transmission of each image independently, and we used the buffer size, β, to control the amount of compression. If $\beta \geq \lambda(x) \times size(x)$, then the entire image, x, fits into the buffer and can be transmitted. However, if $\beta < \lambda(x) \times size(x)$, some bits must be sacrificed. As β decreases, the effective compression factor increases. Given an image, x, we calculate the effective compression ratio, C, of the compressed image x' as

$$C(x') = \frac{size(x)}{size(x') + size(x'.p)}$$

where $x'.p$ is the priority map associated with x'.

We evaluate the quality of the compressed image as the amount of scientific information it contains. Dolinar et al. [3] quantified this by assessing the accuracy of a classifier on the compressed images compared to its performance on the original versions of the same images. In this application, we are concerned with the ability to correctly extract rock features such as albedo and texture. Therefore, we calculate the *fidelity*, F, of a compressed image x' in terms of a feature f_i as

$$F(x', f_i) = 1 - \frac{|x.f_i - x'.f_i|}{x.f_i}.$$

Fidelity is a measure of how well the compressed image preserved the information in feature f_i. A value of 1.0 indicates that f_i was completely preserved. A value of 0.0 means that the calculated value for f_i is 100% off of the true value. Fidelity can go negative, indicating $> 100\%$ error.

Experimental Results. Each image was compressed and transmitted with a priority map that identified rock regions as high priority. For comparison, we also transmitted the same image with a blank priority map. Our goal was to determine whether the priority map produced higher science fidelity in the compressed images.

We used 24 features to characterize each rock. The first two features are the mean and variance of the rock's albedo. The remaining features are 11 average texture features and their variance values. The feature textures are calculated using a Gabor filter at four orientations and three radius values [2]. We omitted one texture mean/variance feature pair because we observed zero values for $x.f_i$, which renders $F(x', f_i)$

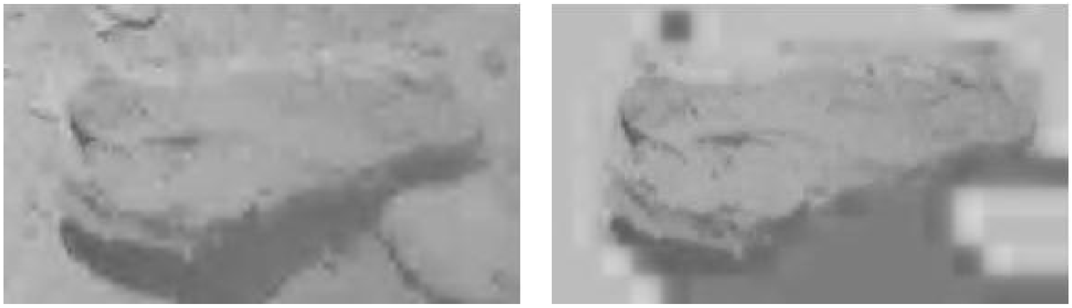

(a) Without priority information (77.6% fidelity) (b) With rock marked high priority (82.6% fidelity)

Figure 1: A FIDO image, compressed 33:1, cropped to 284 × 165 (15% of the image) to show detail.

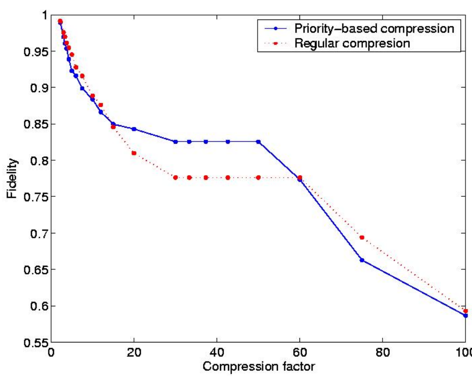

Figure 2: Science fidelity, averaged over 24 features.

undefined. The remaining 24 features provide a sufficient basis for evaluating overall fidelity.

Figure 2 compares the average fidelity obtained from compressed versions of the same images both with and without priority information. The compression ratio ranges from about 2, obtained by lossless compression when no bits need to be discarded, up to 100, corresponding to 0.08 bits per pixel.

We observe that for compression ratios less than about 15, the science-based priority map information has little effect on the fidelity of the compressed image. However, at higher compression levels, we observe clear evidence that the priority maps enable preservation of significantly more detail than regular compression. Figure 1 shows an example of a rock that has been compressed at a ratio of 33 to 1. Figure 1(a) shows the rock when compressed using a uniform priority map, and Figure 1(b) shows the same rock compressed by ROI-ICER with a priority map identifying rock regions as high priority. Detail of the surrounding sand has been sacrificed in favor of preserving rock details.

We also observe that for very high compression levels ($C(x) > 60$) the fidelity of the regular compression method is sometimes higher than that of the priority-based method. We observe this effect because, at very high compression levels, the size of the priority map, $x'.p$, takes up a significant portion of the bit allocation. The priority map is crucial for decompression at the receiving end, so it must be sent using lossless compression. The priority maps for our test images took either one or two packets (1 or 2 kB, respectively). At a compression ratio of 100, each image was compressed down to just three packets total. Therefore, the images produced by ROI-ICER at that compression ratio were, in some cases, permitted only a single packet of image data. In contrast, the uniform priority map used to generate the regular compression results never took more than a single packet. The fidelity obtained at this level of compression is so low that we consider this to be a degenerate case. The compression is so high that the resulting images contain almost no useful information.

Conclusions. Our experiments have shown that, for moderate compression levels, science-based priority maps can significantly improve the observed fidelity of the compressed data. In continuing work, we plan to integrate an automated rock finder that can create the necessary priority maps without manual intervention. We also aim to extend these methods to integrate other science information into the priority maps.

Acknowledgments. We wish to thank several co-workers for manually labeling the rocks in these images. We specifically thank Dominic Mazzoni for his assistance with the feature extraction. This work was funded by the Interplanetary Network Directorate Program and carried out at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

References

[1] A. Castano, et al. Intensity-based rock detection for acquiring onboard rover science. LPSC, 2004.

[2] R. Castano, T. Mann, and E. Mjolsness. Texture analysis for Mars rover images. In *Applications of Digital Image Processing XXII, Proceedings of SPIE*, vol. 3808, 1999.

[3] S. Dolinar, et al. Region-of-interest data compression with prioritized buffer management (III). In *Proceedings of the Third Annual Earth Science Technology Conference*, 2003.

[4] A. Kiely and M. Klimesh. The ICER progressive wavelet image compressor. *The Interplanetary Network Progress Report*, 42-155, 2003.