

# Lossy Image Compression and Stereo Ranging Quality from Mars Rovers

A. Kiely,<sup>1</sup> A. Ansar,<sup>2</sup> A. Castano,<sup>2</sup> M. Klimesh,<sup>1</sup> and J. Maki<sup>3</sup>

*We examine the impact of onboard lossy image compression on the quality of automated ranging results derived on Earth from stereo image pairs acquired from Mars rovers. To assess the quality of stereo ranging results under lossy compression, we acquire stereo image data sets of scenes containing surveyed “ground truth” points. We examine the loss of range estimates and the reduction in ranging accuracy as compressed bit rate decreases. We define an image distortion metric that takes into account the effect of some of the processing steps performed in the automated stereo ranging process. Based on empirical stereo image data, we develop a modification to the ICER image compressor that is designed to optimize this distortion metric rather than mean-square-error distortion. The modification does not increase the complexity of ICER and is shown to improve stereo ranging quality as a function of compressed bit rate on test scenes. The modification does not exploit correlations between left and right images in a stereo pair.*

## I. Introduction

We examine onboard lossy compression of stereo image pairs on Mars rovers. For the Mars Exploration Rover (MER) mission, stereo image pairs from the rovers are compressed onboard as two independent images using the ICER image compressor [5,6]. The compressed images are transmitted to Earth, where automated stereo correlation algorithms are applied to the reconstructed images to assess the terrain surrounding a rover. The resulting data products are used extensively in Mars rover operations for rover navigation and instrument placement.

The Mars Science Laboratory (MSL) mission will also make extensive use of stereo imagery. MSL will include the same engineering cameras as MER—the Hazard Avoidance Cameras (HazCams) and Navigation Cameras (NavCams)—and will use the ICER image compressor for lossy compression of images from these cameras.

---

<sup>1</sup> Communications Architectures and Research Section.

<sup>2</sup> Mobility and Robotic Systems Section.

<sup>3</sup> Space Experiments Systems Section.

The research described in this publication was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

The use of lossy (rather than lossless) image compression significantly increases the amount of imagery returned from Mars over the constrained deep-space communications channel, but the resulting distortion inherent in the reconstructed images reduces stereo correlation performance. For this reason, high data volumes are often required for stereo image pairs used for engineering purposes—for example, typically 8 Mbits of compressed image data (4 bits/pixel)<sup>4</sup> are used for a stereo image pair used for MER robotic arm placement [2].

In this article, we examine the impact of lossy compression on the quality of stereo ranging data products and present a simple modification to the ICER image compressor that improves the trade-off between stereo ranging quality and compressed bit rate.

Improvements in the stereo correlation quality achieved on compressed stereo imagery could be used to reduce the downlink data volume for stereo image pairs, thus making more bandwidth available for science purposes. Moreover, increasing the effective range of stereo ranging data products can significantly reduce planning time during a mission and increase mission capabilities, e.g., by increasing rover driving distance.

We are primarily interested in the quality of the data product produced by the stereo correlation algorithm applied to image pairs, as opposed to a more generic quality metric that might be measured for each image independently. That is, we would be content with an image compression approach that provides accurate stereo ranging information at lower bit rates, even if the images have higher mean-square-error (MSE) distortion or are less appealing visually. We're most interested in stereo image pairs from NavCams and HazCams, since these are most used for navigation, robotic arm placement, etc.

## A. Stereo Ranging Process

The stereo ranging process combines camera calibration, stereo correlation, and triangulation processes. Camera calibration produces a model of the relationship between points in the world and their images on the sensor. Due to optical effects, this relationship is not linear but instead tends to distort the image; extreme examples of this image distortion are common on images from very large angle lenses and fish-eye lenses. The distortion diminishes as the field of view of the lens narrows. The camera calibration procedure models both the intrinsic and the extrinsic parameters of a camera. The intrinsic parameters are those that model the internal characteristics: the focal length, the pixel size, the center of the image, and parameters of some function that describes the non-linear image deformation, including radial distortion. The extrinsic parameters describe the position and orientation of the camera with respect to a reference coordinate frame. At JPL, the most commonly used nonlinear camera models are the CAHVOR and CAHVORE models,<sup>5</sup> consisting of 16 and 20 parameters, respectively, that model both intrinsic and extrinsic parameters [3].

The first step of the ranging process is to linearize and rectify the left and right images acquired by the stereo cameras. The process of linearization consists of removing the radial distortion from an image, producing as output the image that would have been acquired if the camera had no radial distortion—straight lines in the world become straight lines in the linearized image. For this article and for MER operations, the linearized camera model used is the 11-parameter CAHV model [13]. Rectification is the process of aligning the two linearized images so that points in three-dimensional space map to the same scan line in both the left and right cameras.

The next step of the ranging process is to apply stereo correlation algorithms to produce an estimate of the disparity map, which is a record of the change in position between pixels in the left and right

---

<sup>4</sup> Throughout this article, compressed bit rates refer to the ratio of the total number of compressed bits to the *total* number of pixels in a stereo image pair, i.e., left and right cameras each contribute to the number of pixels counted.

<sup>5</sup> The terms CAHVOR, CAHVORE, CAHV are used as shorthand for the set of model parameters included (center, axis, horizontal, vertical, optical, radial, entrance).

images. For example, if the disparity map is referenced to the left image, each value in the map indicates the number of pixels one would have to translate the image of a 3-D point in the left image to find the image of the same point in the right image. In effect, the disparity map is the real output of the ranging process as, once we have the disparity information, we can use the extrinsic parameters of the calibration procedure to triangulate the position of the 3-D point in the world. The effort is in obtaining the best disparity map possible.

An initial estimate of the disparity map is obtained from the JPLstereo program [4], which was designed to run in real time. This program receives two 8-bit linearized and rectified images with their corresponding CAHV models. The rectification guarantees that any 3-D point lies on the same row of both the left and right image, which reduces the search for matches from a 2-D search over the image to a fast 1-D search over an image row. To reduce the effects of offset variations in the sensitivities of the cameras, this search is not performed on the raw images but instead on a high-pass or band-pass filtered version of the images produced by convolution with a difference of box filters to approximate a Laplacian.

The initial disparity map estimate produced by the JPLstereo program serves as a starting point for the Marscor3 program, which produces the final disparity map estimate. Unlike JPLstereo, Marscor3 is not designed for real-time performance but rather for producing a more accurate and dense stereo map. Marscor3 does not rely on the one-dimensional constraint but instead performs a two-dimensional search over a region; the disparity map from JPLstereo is used as an initial guess to start the search. Furthermore, this search is done with an affine transformation that allows the two images to differ slightly due to viewpoint variation. Marscor3 also uses heuristics to fill in some of the holes in the disparity map. For MER operations and for our experiments, a  $3 \times 3$  low-pass box filter is applied to the linearized and rectified left and right images input to Marscor3, in an effort to reduce the impact of compression-induced image distortion [2].

A given pixel in one image may not be successfully matched with a corresponding pixel in the other image. Mainly, three things may happen: (1) the 3-D point seen in one of the images may be occluded from view in the other image, in which case, no match exists, (2) the 3-D point may be visible in both images but located in an area without texture and thus not allowing a strong match, in which case the program will conservatively avoid estimating a match, or (3) the 3-D point seen from one image may be incorrectly matched to the image of a different 3-D point, giving rise to a mismatch. Failing to find a match gives rise to a “hole” in the estimated range, while a mismatch yields an incorrect range value.

A range image is produced by combining the disparity map and the extrinsic parameters of the calibration procedure. For each pixel in the linearized and rectified left or right image, we have either a range estimate (i.e., an estimate of the distance from the camera to the point in the scene corresponding to the pixel) or a hole. For the purposes of the discussion here, we’ll view this range image as our final product of interest.

## B. Stereo Cameras

The following stereo cameras are of interest for this article:

- The MER HazCam [7], which has a focal length of 5.58 mm, a 124-deg  $\times$  124-deg field of view, and a 10-cm stereo baseline.
- The MER NavCam, which has a focal length of 14.67 mm, a 45-deg  $\times$  45-deg field of view, and a 20-cm stereo baseline.
- The MER Panoramic Camera (PanCam), which has a 43-mm focal length, a 16-deg  $\times$  16-deg field of view, and a 30-cm baseline.

- Our narrow field-of-view camera (NFOV-Cam), which consists of two Point Grey Research Scorpion 20SO cameras separated by a 20-cm baseline and using 16-mm Pentax lenses. The images are cropped to  $1024 \times 1024$  pixels to match MER cameras, resulting in a 16-deg  $\times$  16-deg field of view. The NFOV-Cam, which is similar to PanCam, is an inexpensive portable camera that allowed us to collect good stereo test imagery without being dependent on the availability of experimental rovers.

The MER cameras all produce  $1024 \times 1024$ -pixel images with a bit depth of 12 bits/pixel.

### C. Related Research

Nagura [8] has studied onboard lossy compression of Earth orbital stereo imagery used for height detection. The system he considered makes use of three or four views of a scene to provide good height resolution while reducing data loss due to occlusions. Nagura proposes one-dimensional motion compensation to provide a coarse description of the difference between images. Locally averaged motion vector images are losslessly compressed using the JPEG-LS compressor [12]. Lossy JPEG2000 compression [11] is applied to a reference image and the images of the difference between an image and the estimate of that image based on the motion vector information. Note that the multi-view system produces more difference images than a conventional stereo application, and so effective compression of these difference images is likely to have a larger impact on overall performance.

To obtain a performance metric to assess stereo compression effectiveness, in [8–10] Nagura evaluates the fraction of “error pixels” in a scene—that is, the fraction of pixels whose height estimate is in error by more than some threshold. We know of no other cases in the literature of metrics devised to indicate ranging quality from compressed stereo imagery, and in fact the work of Nagura is the only we know of that specifically aims at onboard lossy compression of stereo imagery.

## II. Stereo Data Sets with Ground Truth

The automated stereo ranging process is subject to errors even when it is applied to the original images. We would like to assess the quality of stereo ranging data products both with and without lossy data compression so that we can better gauge the relative quality reduction due to lossy compression. In some situations, the additional error introduced by lossy compression may turn out to be small compared to the limitations inherent in the automated ranging process.

### A. Acquisition

To serve as test data sets, we acquired two scenes each from a HazCam<sup>6</sup> and our NFOV-Cam. For each scene, we collected a stereo image pair along with several “ground truth” test points whose positions were measured using a Leica total station, which is a high-precision laser surveying tool. Figures 1 and 2 show the left images from each scene. Figures 3 and 4 show the locations of the surveyed points in each scene along with the reprojections described below.

For both NFOV-Cam scenes, ground truth points could be more reliably located in the right images because they were in better focus than the left images. Consequently, stereo ranging operations were performed using the right image as a reference for NFOV-Cam scenes. The left camera was used as the reference for the HazCam scenes, which is the more common practice.

A camera calibration procedure was performed as part of the data collection. The calibration procedure is designed to infer the intrinsic and extrinsic parameters of a camera from a combination of image data

---

<sup>6</sup>The HazCam scenes were acquired using the Surface System Testbed (SSTB) rover, whose cameras are identical to the MER cameras.

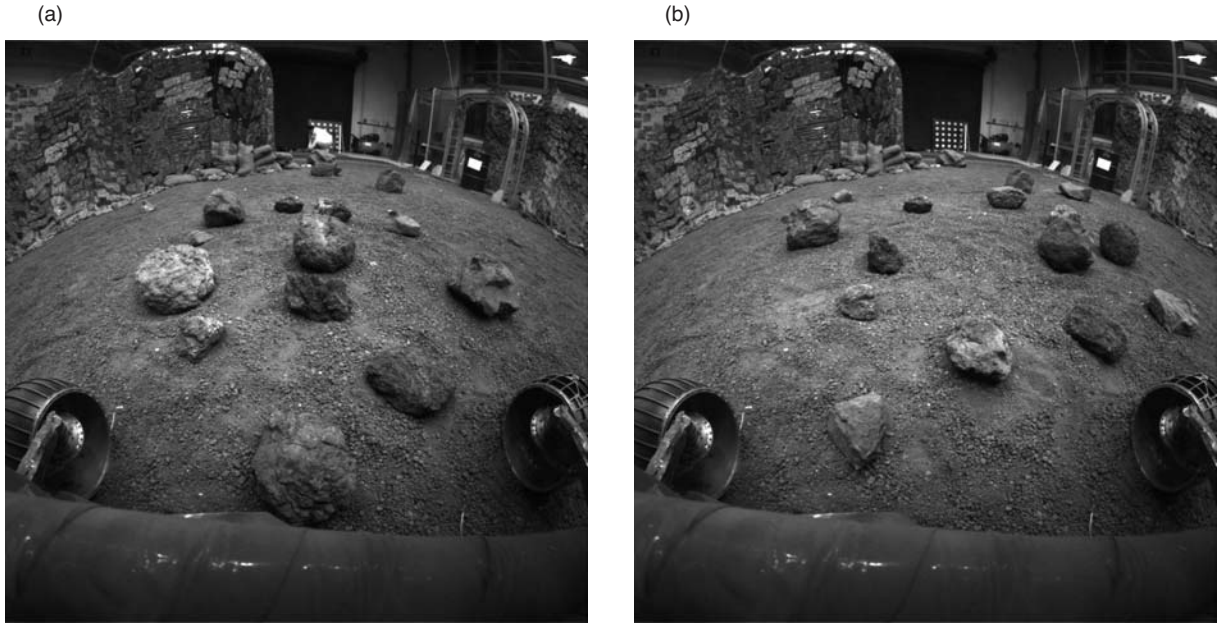


Fig. 1. Left images from HazCam: (a) scene A and (b) scene B.

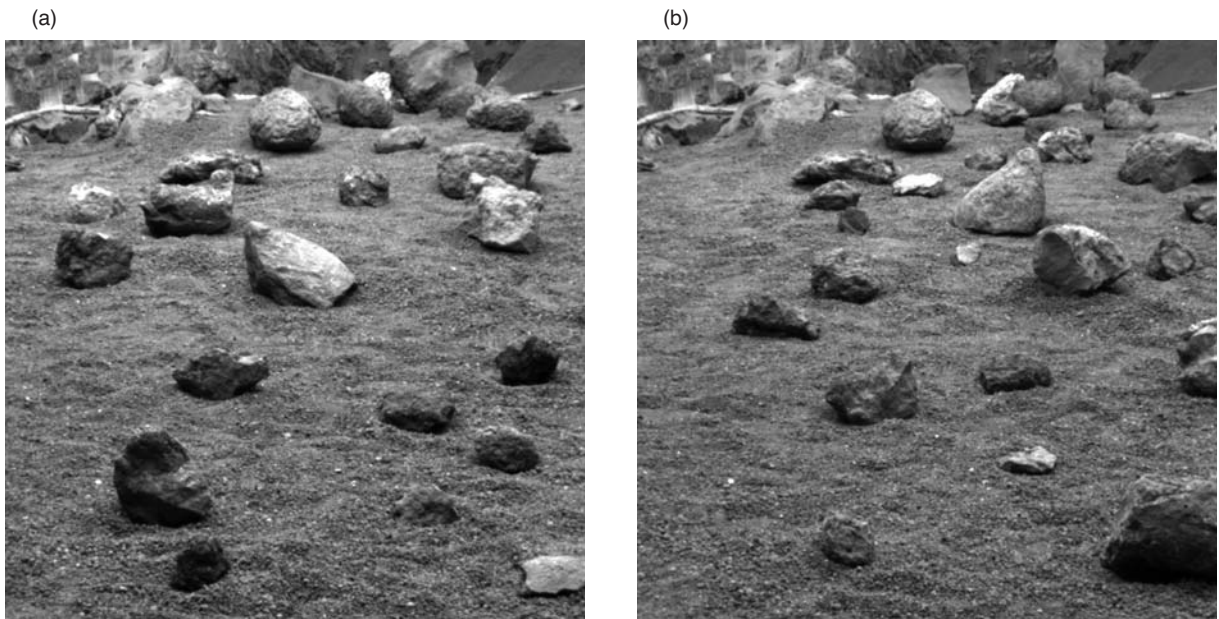
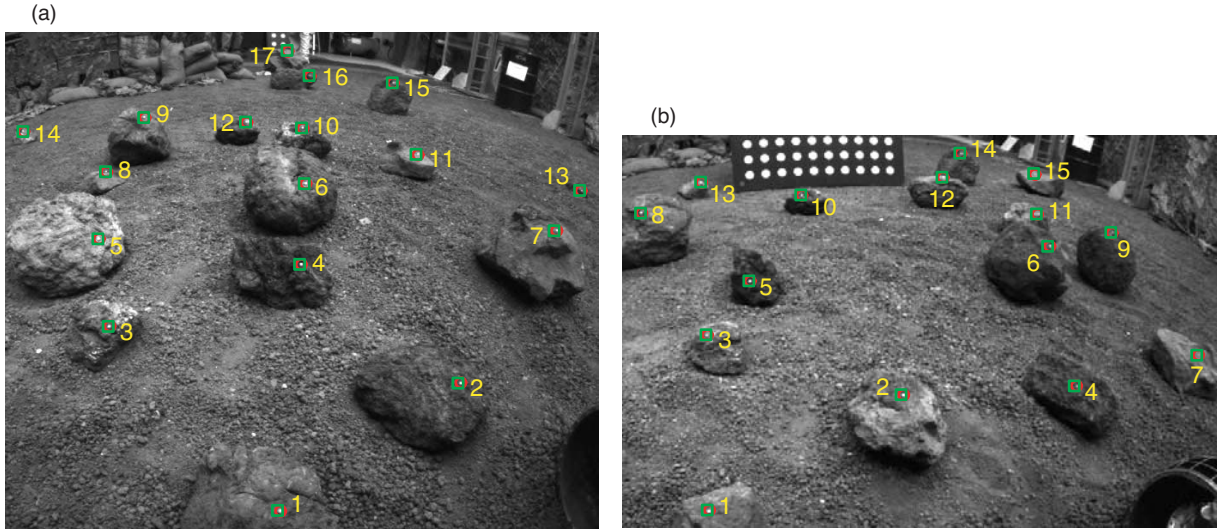
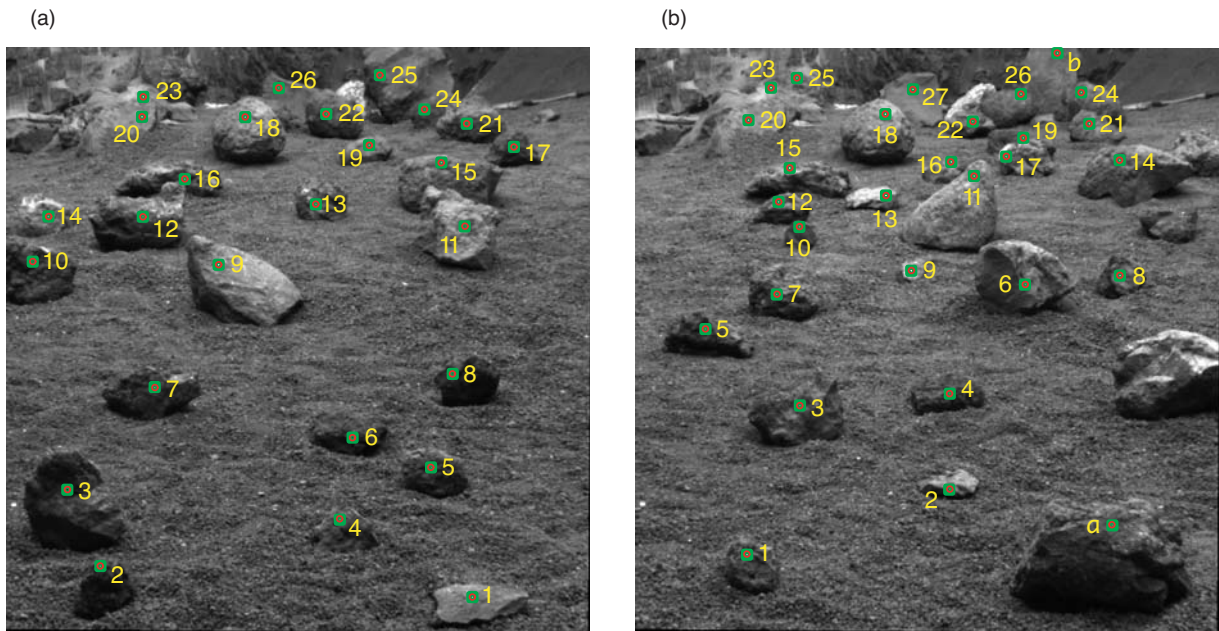


Fig. 2. Left images from NFOV-Cam: (a) scene A and (b) scene B.

and 3-D measurements. For our purposes, calibration data are also used to determine range values of test points (by determining the camera position). The procedure used for results in this article is similar to that used to calibrate the MER cameras. A 1-m  $\times$  1-m planar target with a 10  $\times$  10 grid of dots [a portion of which appears in Fig. 3(b)] is imaged in various positions to cover both the full field of view of the stereo pair and, to the extent possible, the range of distances spanning the work environment of the stereo cameras. At each position, the 3-D locations of the four corners of the grid are measured using the total station. All measurements are in the fixed coordinate frame of the total station. Since the target



**Fig. 3. Details from left images of HazCam: (a) scene A and (b) scene B, including clay markers. Red circles are centered around visually located points, and green boxes are centered at the reprojection of each surveyed point.**



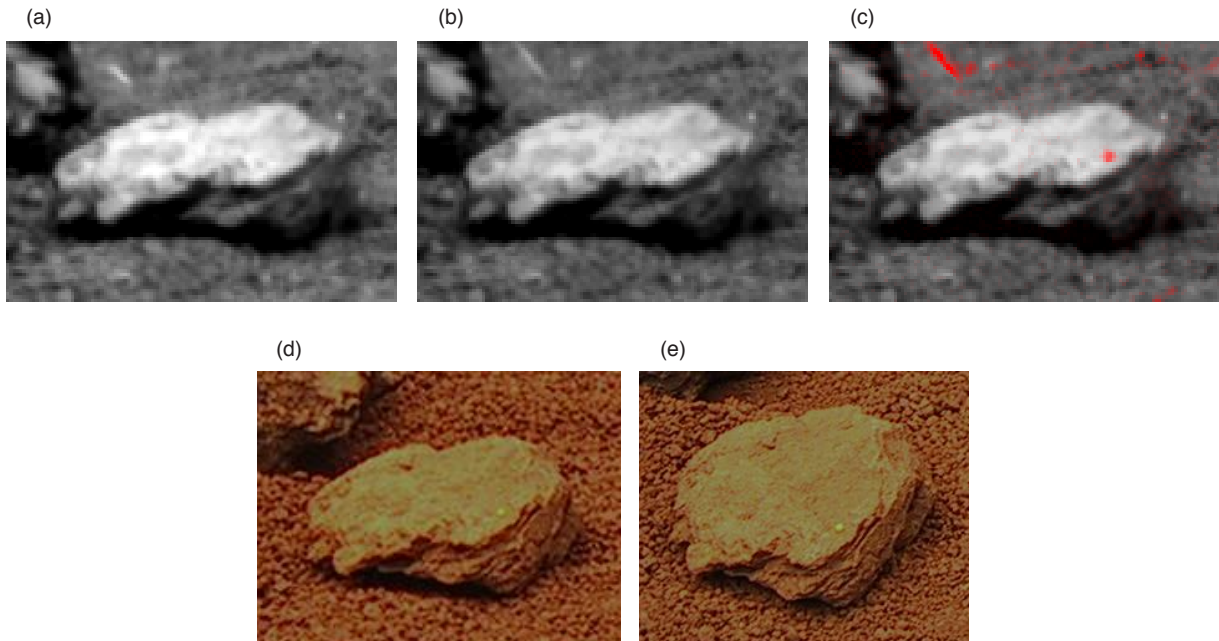
**Fig. 4. Linearized and rectified right images from NFOV-Cam: (a) scene A and (b) scene B, including clay markers. Red circles are centered around visually located points, and green boxes are centered at the reprojections of the surveyed points.**

grid is uniform, measuring the four corners gives the 3-D location of each dot in the grid via interpolation. Image-processing software extracts the pixel coordinates of the center of each dot on each target. Thus, we have the 3-D coordinates of each dot on each target in a fixed frame as well as the corresponding pixel coordinates for each camera. Calibration then amounts to a least-squares optimization over the parameters of the CAHVOR or CAHVORE model to produce the best match between 3-D data and projected 2-D data.

Because the calibration process is somewhat time consuming, the camera was kept in a fixed position between scenes so that a single calibration procedure could be used for both scenes from each camera. The large rocks were moved and the gravel was trampled somewhat between scenes, but a significant portion of the two scenes for each camera is identical, and so one should keep in mind that the amount of diversity in our data is less than we would have if we had acquired two completely different scenes for each camera.

Several small clay markers were placed in each scene to mark the locations of ground truth points. The location of each clay marker was surveyed using the total station, and the coordinates were transformed into the camera frame of reference so that a range value could be calculated for each point. Each scene was acquired by the stereo camera both with (Figs. 3 and 4) and without (Figs. 1 and 2) the clay markers. Locating the markers in the images is relatively straightforward using images of the difference between the scene with and without the markers. The positions of markers in the difference images tended to be conspicuous because the scenes were acquired indoors under unchanging lighting conditions and because of the contrast between the bright marker and darker background. Color photographs from a handheld camera were also used to help confirm the locations of markers in the scene. Figure 5 illustrates how the various images can be used to locate a marker in the scene. As a check, we also reprojected the surveyed points into the linearized and rectified left or right image of each scene. Figures 3 and 4 show that the two methods agree reasonably well. Stereo range estimation is performed on scenes without markers; otherwise, the high contrast between the marker and background might unfairly increase the ranging accuracy at each marker. Figures 6 and 7 show color-coded range images.

For a significant number of the test points, we placed markers on portions of rocks that (locally, over a few square centimeters) were approximately orthogonal to the camera’s line of sight. The motivation for this placement is to increase the likelihood that the test point is in a cluster of pixels with similar range estimates, so that an error of a few pixels in locating the marker center would not tend to cause large differences in the stereo-derived range estimate for the marker.



**Fig. 5. Details of right linearized and rectified images of NFOV-Cam scene B in the vicinity of point B13: (a) scene without clay marker, (b) scene with clay marker, and (c) using a red tint to highlight pixels that are brighter in (b) than in (a). The location of the clay marker is difficult to determine in (b) but conspicuous in (c). The color photographs (d) and (e) from the handheld camera help to confirm the marker location.**

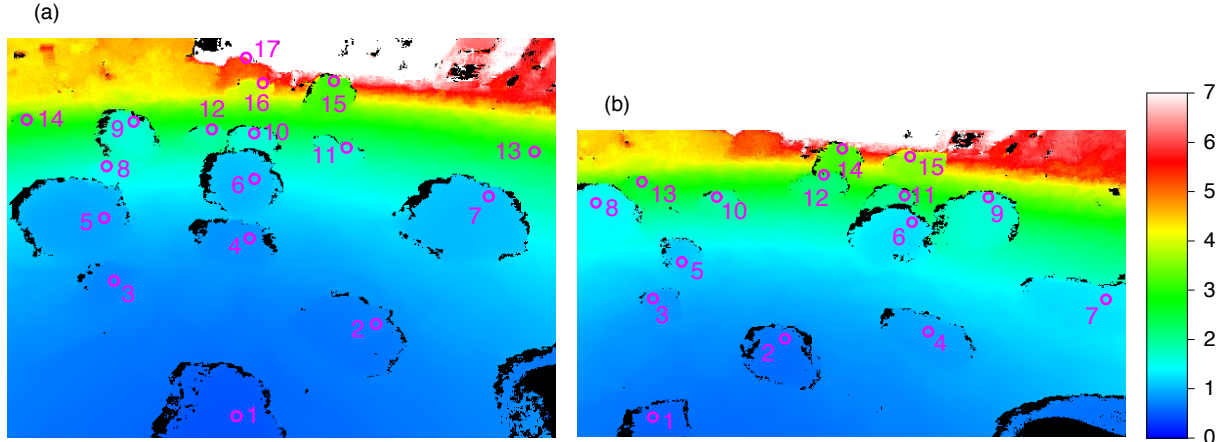


Fig. 6. Cropped range results from HazCam: (a) scene A and (b) scene B. The circles are centered around visually located test points. The scale at right indicates range in meters. Pixels shown in black have no range estimate, and pixels in white are estimated at 7 meters or more.

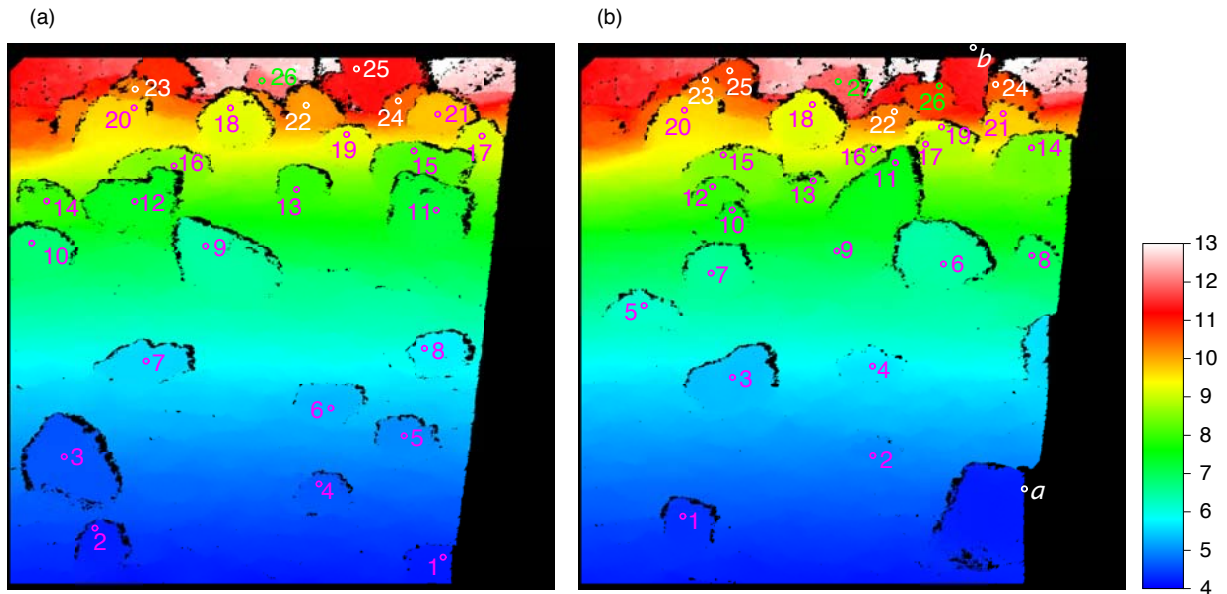
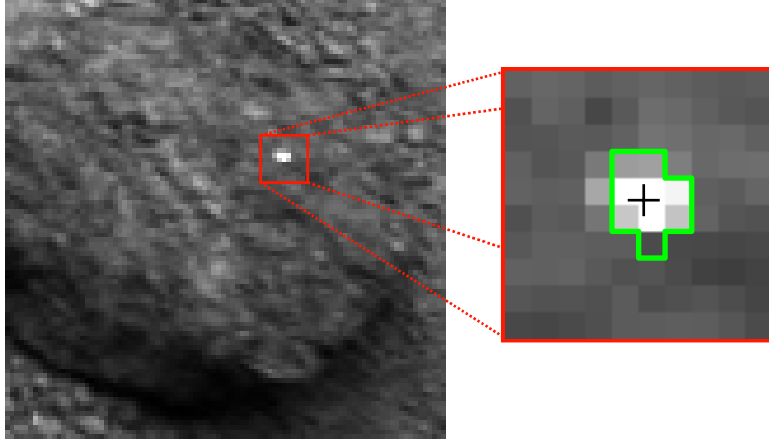


Fig. 7. Range results from NFOV-Cam: (a) scene A and (b) scene B. The circles are centered around visually located test points. The scale at right indicates range in meters. Pixels shown in black have no range estimate, and pixels in white have range estimated at 13 meters or more.

A point was assigned at the approximate center of each marker in the linearized and rectified left or right image by visual inspection. The markers are somewhat larger than a single pixel, and so selecting a precise center for each marker in the photograph is a somewhat subjective and imprecise process. To account for this, for each test point we take the average range estimate over a cluster of 9 pixels, selected as those with centers closest to the nominal (visually located) marker location. Figure 8 illustrates such a cluster.





**Fig. 8.** Detail from left linearized and rectified image of HazCam scene A near test point A2. Black crosshairs indicate the nominal location of the test point. The range estimate for the test point is taken as the average of the 9-pixel cluster outlined in green.

## B. Range Results

Tables 1 through 4 give a comparison of surveyed and stereo-derived range values for the test points, depicted graphically in Fig. 9. For each 9-pixel cluster of stereo-derived range estimates  $x_1, \dots, x_9$ , we tabulate the mean  $\bar{x}$  and sample standard deviation, which for a general number of pixels  $N$  is

$$s = \sqrt{\frac{1}{N-1} \sum_{i=1}^N (x_i - \bar{x})^2}$$

The tables also contain the discrepancy between  $\bar{x}$  and the surveyed range as well as an estimate  $\sigma$  of the uncertainty in the stereo-derived range. For a point at range  $r$ , an uncertainty of  $\Delta d$  pixels in the inter-image disparity translates to a range uncertainty of [1]

$$\sigma = \frac{r^2}{FB} \Delta d$$

Here  $F$  is the camera focal length in pixels and  $B$  is the baseline distance. For the HazCams,  $F = 278.1$  and  $B = 0.1$  meters [7]. For the NFOV-Cams, the camera calibration process yields  $F = 3702.4$  and  $B = 0.19978$  meters. For  $r$ , we use the range value obtained from surveyed data. For expected sub-pixel accuracy, we use  $\Delta d = 1/3$ , which is the rule of thumb used by MER for sub-pixel accuracy in correlation.

The range of a test point is the distance from the point to the camera's origin (i.e., the origin of the camera's coordinate frame of reference). The physical location of the camera's origin is inside the camera and thus cannot be surveyed. So, to compute range estimates of test points from surveyed data, we use the camera calibration information to map the location of the camera's origin into the coordinate system used by the surveying tool.

Based on the total station manufacturer's specifications, the standard deviation of the error in the surveyed coordinates of each test point is 3 mm. In our experience with camera calibration, using two different calibration data sets for a given camera configuration yields camera position estimates that differ by much less than 1 mm, and so we conclude that locating the camera origin in the coordinate system

**Table 1. Range data for test points in HazCam scene A.**

Point	Surveyed range, cm	Stereo-estimated range (9-pixel) cluster		Discrepancy, cm	Range uncertainty $\sigma$ , cm
		Mean, cm	Standard deviation $s$ , cm		
A1	44.6	45.1	0.1	0.5	0.2
A2	68.6	69.3	0.2	0.7	0.6
A3	71.5	72.2	0.3	0.7	0.6
A4	80.8	81.1	0.2	0.3	0.8
A5	91.9	92.6	0.4	0.7	1.0
A6	103.6	104.4	0.3	0.8	1.3
A7	105.7	107.0	0.4	1.3	1.3
A8	148.0	149.7	2.1	1.7	2.6
A9	163.0	165.3	2.4	2.4	3.2
A10	164.2	164.0	0.8	-0.2	3.2
A11	170.0	172.8	1.0	2.8	3.5
A12	192.7	196.5	0.6	3.8	4.5
A13	245.1	251.8	1.3	6.7	7.2
A14	270.7	275.8	2.4	5.1	8.8
A15	316.6	320.1	1.7	3.6	12.0
A16	395.4	427.2	12.7	31.9	18.7
A17	523.3	702.9	75.3	179.7	32.8

**Table 2. Range data for test points in HazCam scene B.**

Point	Surveyed range, cm	Stereo-estimated range (9-pixel) cluster		Discrepancy, cm	Range uncertainty $\sigma$ , cm
		Mean, cm	Standard deviation $s$ , cm		
B1	53.5	53.8	0.1	0.3	0.3
B2	62.7	63.0	0.2	0.3	0.5
B3	85.1	85.2	0.4	0.2	0.9
B4	88.5	89.1	0.4	0.6	0.9
B5	106.5	107.3	0.3	0.8	1.4
B6	125.6	126.3	0.3	0.6	1.9
B7	130.2	130.7	0.3	0.6	2.0
B8	141.5	142.3	0.8	0.7	2.4
B9	167.5	168.4	0.9	0.9	3.4
B10	192.0	191.9	0.5	-0.1	4.4
B11	204.2	210.5	3.2	6.2	5.0
B12	219.5	229.7	10.6	10.2	5.8
B13	250.8	249.6	2.0	-1.2	7.5
B14	317.1	318.7	1.7	1.7	12.0
B15	363.4	366.2	6.4	2.9	15.8

**Table 3. Range data for test points in NFOV-Cam scene A.**

Point	Surveyed range, cm	Stereo-estimated range (9-pixel cluster)		Discrepancy, cm	Range uncertainty $\sigma$ , cm
		Mean, cm	Standard deviation $s$ , cm		
A1	419	420	0.2	0.7	0.8
A2	422	422	0.2	0.0	0.8
A3	460	459	0.2	-0.5	1.0
A4	461	461	0.2	-0.5	1.0
A5	498	497	0.4	-0.7	1.1
A6	527	527	0.2	0.2	1.3
A7	554	554	0.3	0.4	1.4
A8	571	570	0.7	-0.9	1.5
A9	647	646	0.4	-1.0	1.9
A10	681	679	1.3	-2.2	2.1
A11	725	722	0.3	-3.0	2.4
A12	739	739	1.3	-0.3	2.5
A13	792	792	0.8	0.2	2.8
A14	796	796	1.6	0.2	2.9
A15	812	810	1.1	-2.2	3.0
A16	844	843	1.2	-1.0	3.2
A17	913	914	1.2	0.9	3.8
A18	922	922	1.2	0.0	3.8
A19	925	922	0.8	-2.2	3.9
A20	966	964	1.3	-1.6	4.2
A21	980	981	0.8	0.6	4.3
A22	1009	1009	1.0	-0.6	4.6
A23	1027	1030	2.0	3.1	4.8
A24	1050	1049	1.7	-1.7	5.0
A25	1132	1139	2.1	6.8	5.8
A26	1200	1205	0.9	4.7	6.5

used by the surveying tool introduces less than 1-mm additional error. These error sources are smaller than the roughly 5-mm diameter of the clay markers used to locate the test points. We conclude that the range estimates derived from surveyed data have less than 1-cm error.

We note from Fig. 9 that, in the HazCam scenes, the stereo-derived range estimates tend to be higher than range estimates from surveyed data. We suspect that this apparent bias is due to a slight error in the baseline distance value used in the camera models.

### C. Outlier Points

For a large majority of the test points, agreement between surveyed and stereo-estimated range is quite good, with discrepancy typically less than  $\sigma$ . There are six outlier points (HazCam A1, A16, A17, B12, and NFOV-Cam B13, B19) that merit some discussion.

Four of the six outlier points (HazCam A16, A17, B12, and NFOV-Cam B19) are close to the edge of a rock, which presents more of a challenge for stereo correlation algorithms. In addition, the 9-pixel

**Table 4. Range data for test points in NFOV-Cam scene B.**

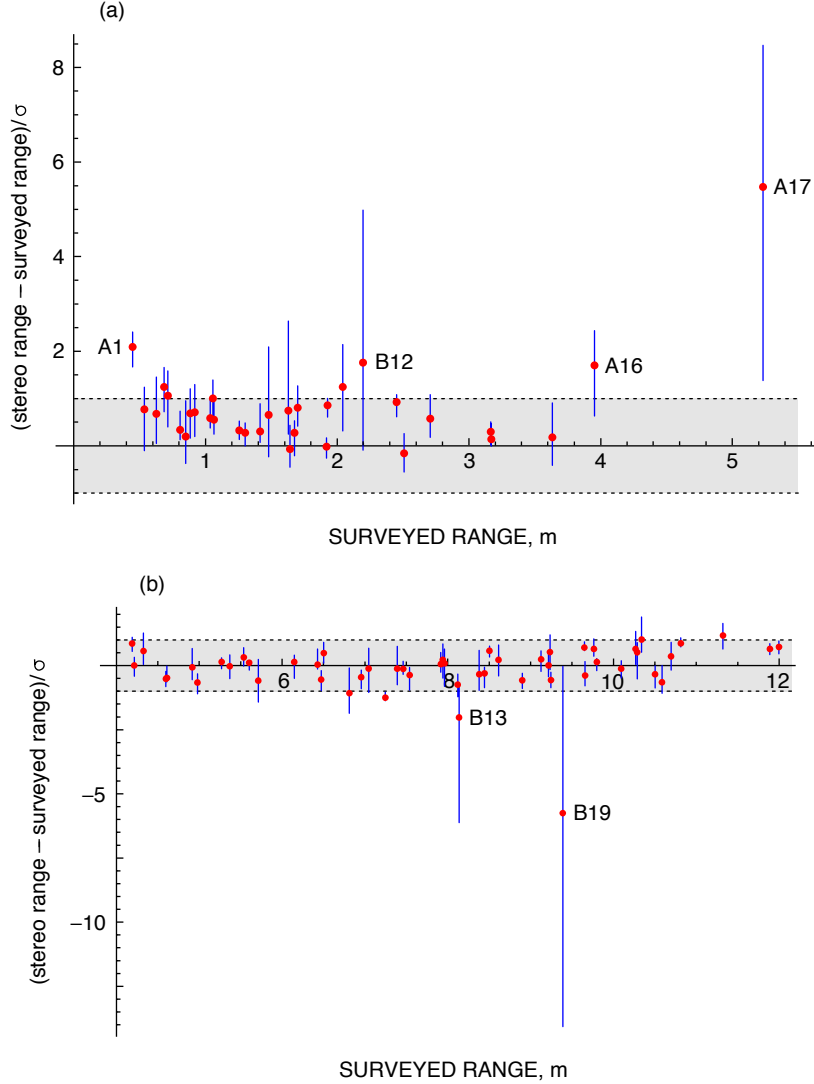
Point	Surveyed range, cm	Stereo-estimated range (9-pixel cluster)		Discrepancy, cm	Range uncertainty $\sigma$ , cm
		Mean, cm	Standard deviation $s$ , cm		
B1	433	433	0.3	0.5	0.8
B2	492	492	0.5	-0.1	1.1
B3	537	537	0.4	0.0	1.3
B4	560	560	0.2	0.2	1.4
B5	615	615	0.5	0.2	1.7
B6	643	643	0.5	0.1	1.9
B7	650	651	0.6	0.9	1.9
B8	696	695	0.5	-1.0	2.2
B9	705	704	1.2	-0.3	2.2
B10	746	746	0.4	-0.3	2.5
B11	754	753	0.7	-0.9	2.6
B12	794	795	1.5	0.6	2.8
B13	814	808	5.7	-6.0	3.0
B14	838	837	1.5	-1.1	3.2
B15	850	852	0.6	1.9	3.3
B16	861	862	1.3	0.7	3.3
B17	890	888	0.7	-2.0	3.6
B18	923	925	1.8	2.0	3.8
B19	939	916	23.6	-22.8	4.0
B20	965	968	0.5	2.9	4.2
B21	976	979	0.9	2.8	4.3
B22	1029	1031	2.5	2.5	4.8
B23	1034	1039	1.9	4.9	4.8
B24	1059	1055	2.3	-3.3	5.1
B25	1070	1071	1.5	1.9	5.2
B26	1081	1086	0.6	4.6	5.3
B27	1189	1193	0.9	4.1	6.4

cluster for such a test point is more likely to include pixels that are off of the rock and thus higher than the true range of the test point. We note, however, that HazCam points A8, A15, and B13 are close to rock edges and don't seem to suffer from the same problem.

NFOV-Cam point B13 is located on a rock with a bright surface, and the left image of the test point is located in a relatively large region of saturated pixels, as illustrated in Fig. 10. It is therefore not surprising that the range accuracy is low in this region.

The final outlier, HazCam point A1, is the closest test point to the camera. The discrepancy between surveyed and stereo-derived range estimates for the point is half of a centimeter, which is comparable to the accuracy of our range estimates derived from surveyed values.

In NFOV-Cam scene B, points *a* and *b* (see Figs. 4 and 7) are near the edge of one or both images, and no range estimates are produced.

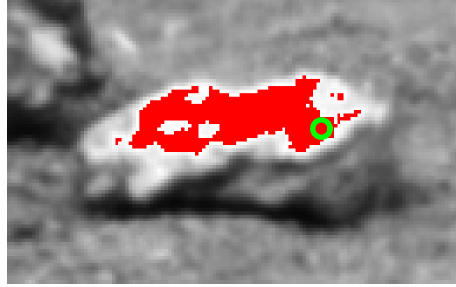


**Fig. 9.** Comparison of surveyed and stereo-derived range estimates for test points in (a) HazCam and (b) NFOV-Cam scenes. Each red point indicates the mean stereo-derived range estimate over the 9-pixel cluster. The blue bars stretch from minimum to maximum range estimate over each 9-pixel cluster. The shaded gray region indicates the interval  $[-\sigma, \sigma]$ .

### III. Test Images

In addition to the stereo images described in Section II, we also make use of images of Mars from the MER rover Opportunity that were originally sent using lossless compression. Table 5 lists the images used in this article.

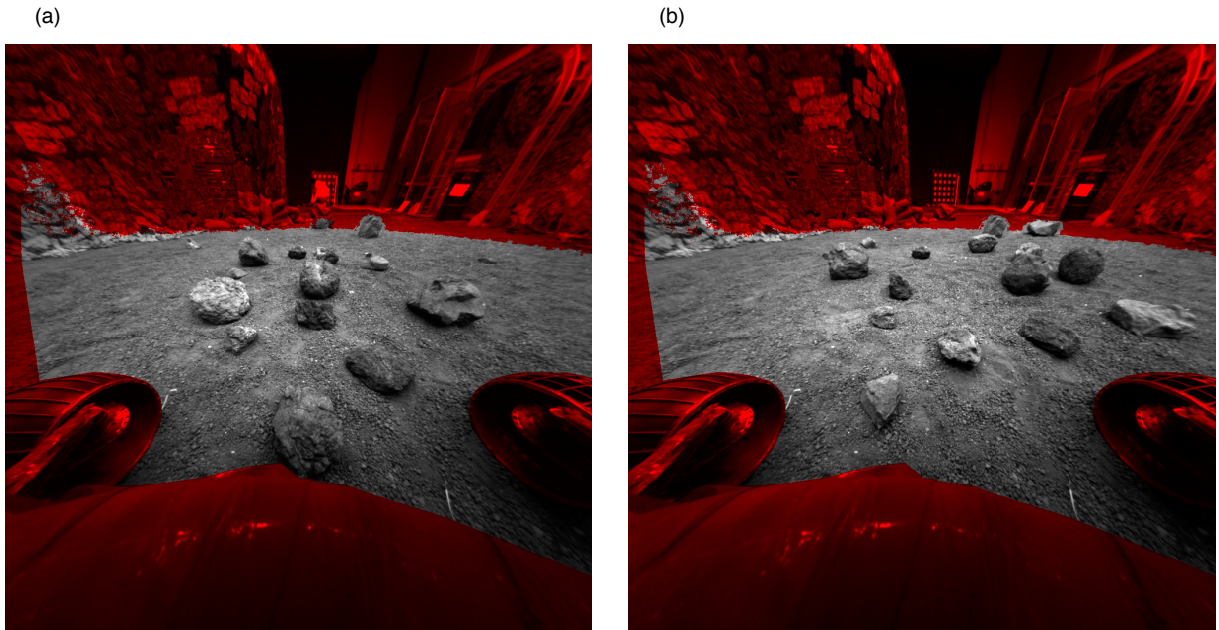
A large fraction of the pixels in HazCam images corresponds to so-called far-field terrain that is beyond the typical range of interest for HazCam stereo imaging (this far-field terrain is more accurately imaged by the Navcams), or corresponds to locations on the rover itself, where range information is less important. In an effort to provide a more relevant assessment of stereo quality, for each left linearized and rectified HazCam image we have defined masks to isolate *high interest* pixels, comprising about 37 percent of the pixels in HazCam scenes A, B, and  $\Omega$ . The masks are illustrated in Fig. 11 for HazCam scenes A and B.



**Fig. 10.** Detail from left linearized and rectified image of NFOV-Cam scene B near point B13, with saturated pixels colored red. The green circle indicates the approximate location of point B13.

**Table 5. Summary of test images.**

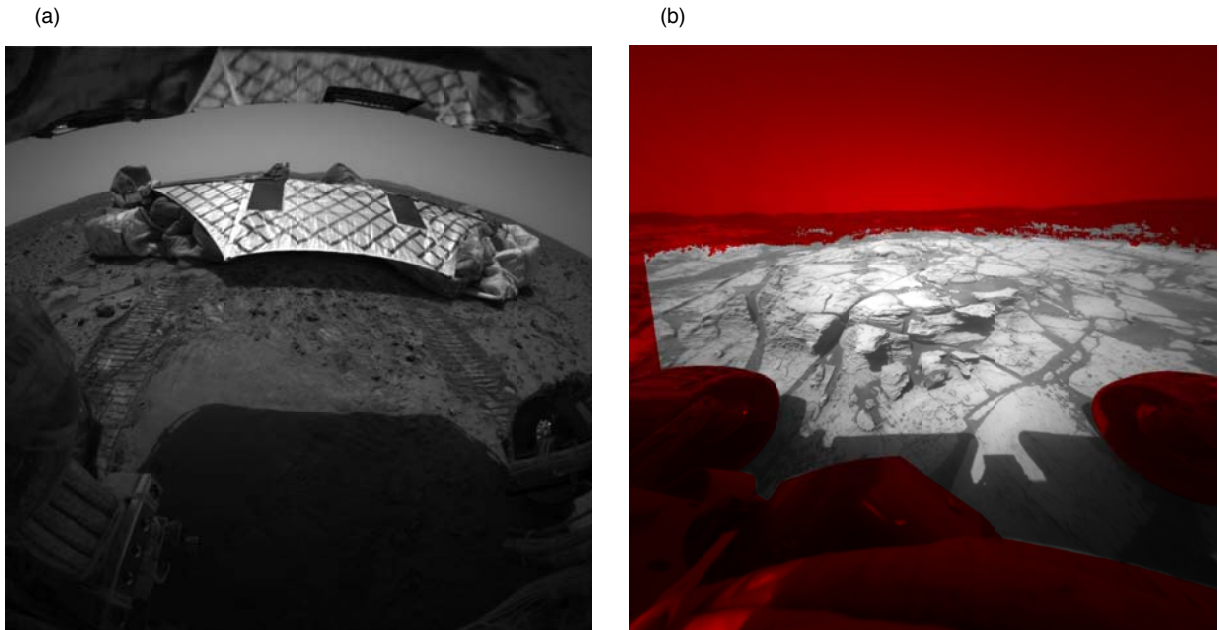
Camera	Label	Source
HazCam	A	See Section II
HazCam	B	See Section II
NFOV-Cam	A	See Section II
NFOV-Cam	B	See Section II
HazCam	$\Psi$	MER 1R178787408EFF5927P1319R0M1 (single image)
HazCam	$\Omega$	MER 1F187750538EFF64KCP1219 (stereo pair)
NavCam	$\Phi$	MER 1N184911329EFF6400P1939R0M1 (single image)
NavCam	$\Theta$	MER 1N138304998EFF2700P1993 (stereo pair)



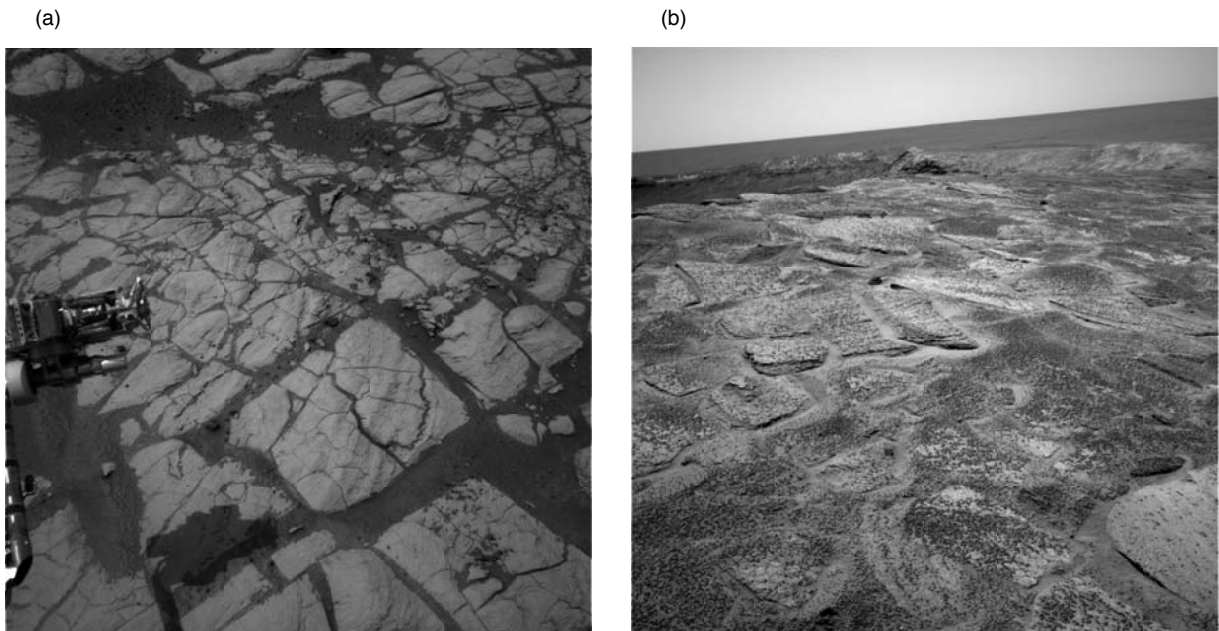
**Fig. 11.** Left linearized and rectified images from HazCam: (a) scene A and (b) scene B, with masked (low interest) pixels shaded red.

The masks exclude pixels located on the rover, pixels more distant than 4 m (based on estimates from the ranging algorithm applied to the original images), and pixels near the left edge of the image that are not visible in the right image.

Figure 12 shows MER HazCam images along with the mask used to identify high-interest pixels in scene  $\Omega$ . Figure 13 shows MER NavCam images.



**Fig. 12. MER HazCam images: (a) image  $\Psi$  and (b) left linearized/rectified image from scene  $\Omega$  with masked (low interest) pixels shaded red.**



**Fig. 13. MER NavCam images: (a) image  $\Phi$  and (b) left image from scene  $\Theta$ .**

## IV. Effect of Lossy Compression on Stereo Range Estimates

### A. General Discussion

For a given compression approach, we'd like to assess the impact of lossy compression on stereo ranging performance as a function of compressed bit rate. Conventional image distortion metrics, such as MSE distortion between the original and reconstructed image, are not necessarily good indicators of stereo ranging quality, and so we'd like to have a more suitable metric to compare different compression approaches.

Whenever lossy image compression is used, the potential effects on the range result for a given pixel are

- (1) The value of the range estimate might change.
- (2) A range estimate might be *lost*, i.e., the correlator might produce a range estimate for a pixel on the original images but fail to produce one under a given lossy compression scenario.
- (3) A range estimate might be *found*, i.e., the correlator might produce a range estimate for a pixel under a given lossy compression scenario but not for the original image pair.

Figure 14 illustrates these three different effects when lossy ICER image compression is applied to HazCam scene A and NFOV-Cam scene A. For the NFOV-Cam scene, the dominant effect of lossy compression at 0.5 bits/pixel seems to be the presence of a few clusters of lost pixels (shown in white). For the HazCam scene, lossy compression to 1 bit/pixel results in some noticeable changes in range estimates (shown in yellow/green) for more distant pixels, a few lost near-field pixels (white), and a noticeable cluster of found pixels (red) on the left rover wheel. The found pixels might represent useful data—the smoothing occurring from the lossy image compression process might better enable the correlator to find matches in this part of the scene.

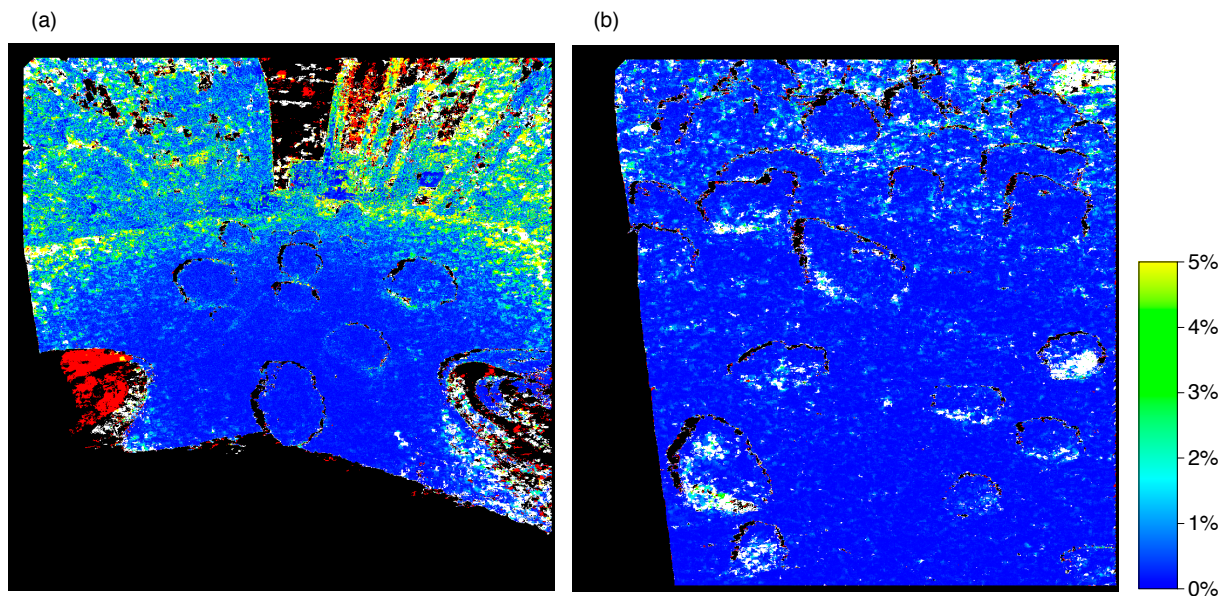


Fig. 14. Illustration of the differences between range estimates produced at lossless compression and using ICER lossy compression at (a) 1 bit/pixel for HazCam scene A and (b) 0.5 bits/pixel for NFOV-Cam scene A. For pixels shown in black, no range estimates were produced under lossless or lossy compression. White pixels indicate lost pixels; red pixels indicate found pixels. The remaining pixels are colored to indicate the absolute value of the percentage difference in range estimate obtained under lossy and lossless compression, using the color scale at right, with yellow pixels indicating a discrepancy of 5% or more.



## B. Effect of Lossy Compression on Test Data Sets

To present an example of the impact of lossy compression on stereo ranging quality, we evaluate stereo ranging performance when the ICER compressor is used on HazCam scenes A and B and NFOV-Cam scenes A and B. In each case, left and right images are compressed to the same bit rate, which is the procedure used for stereo imagery onboard MER.

For each test point, Fig. 15 shows the discrepancy between stereo-derived range estimates (taken as the average over 9-pixel clusters) and surveyed range values, normalized by  $\sigma$ , as we vary the compressed bit rate. Figure 16 shows the fraction of lost and found pixels as a function of compressed bit rate for the HazCam and NFOV-Cam scenes.

We see from Fig. 15 that the range estimates for the test points change as we decrease the compressed bit rate, but the change is generally small compared to the apparent range error obtained under lossless compression, becoming significant only at low bit rates (less than 0.5 bits/pixel for the HazCam scenes and NFOV-Cam scene B, and less than 1.5 bits/pixel for NFOV-Cam scene A). An exception is NFOV-Cam test point B27, which is located on a relatively smooth rock face, and whose range accuracy degrades more noticeably at low bit rates, even failing to produce a range estimate at 0.25 bits/pixel.

As the compressed bit rate changes, outlier points HazCam A17, B12, and NFOV-cam B19 (all located near the edge of rocks) exhibit more fluctuations in stereo-derived range estimates than the remaining test points. Even for these points, however, the apparent range error present when the original images are used seems to be more significant than the additional range error due to lossy compression.

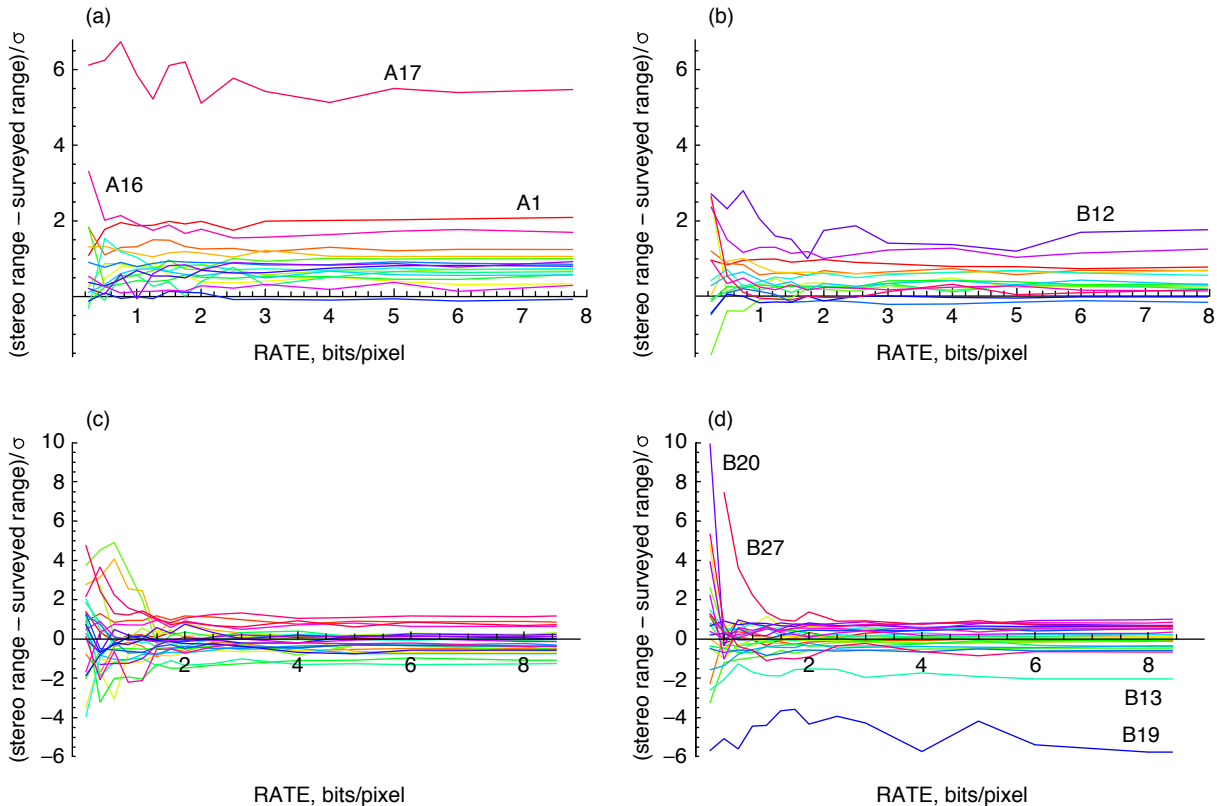
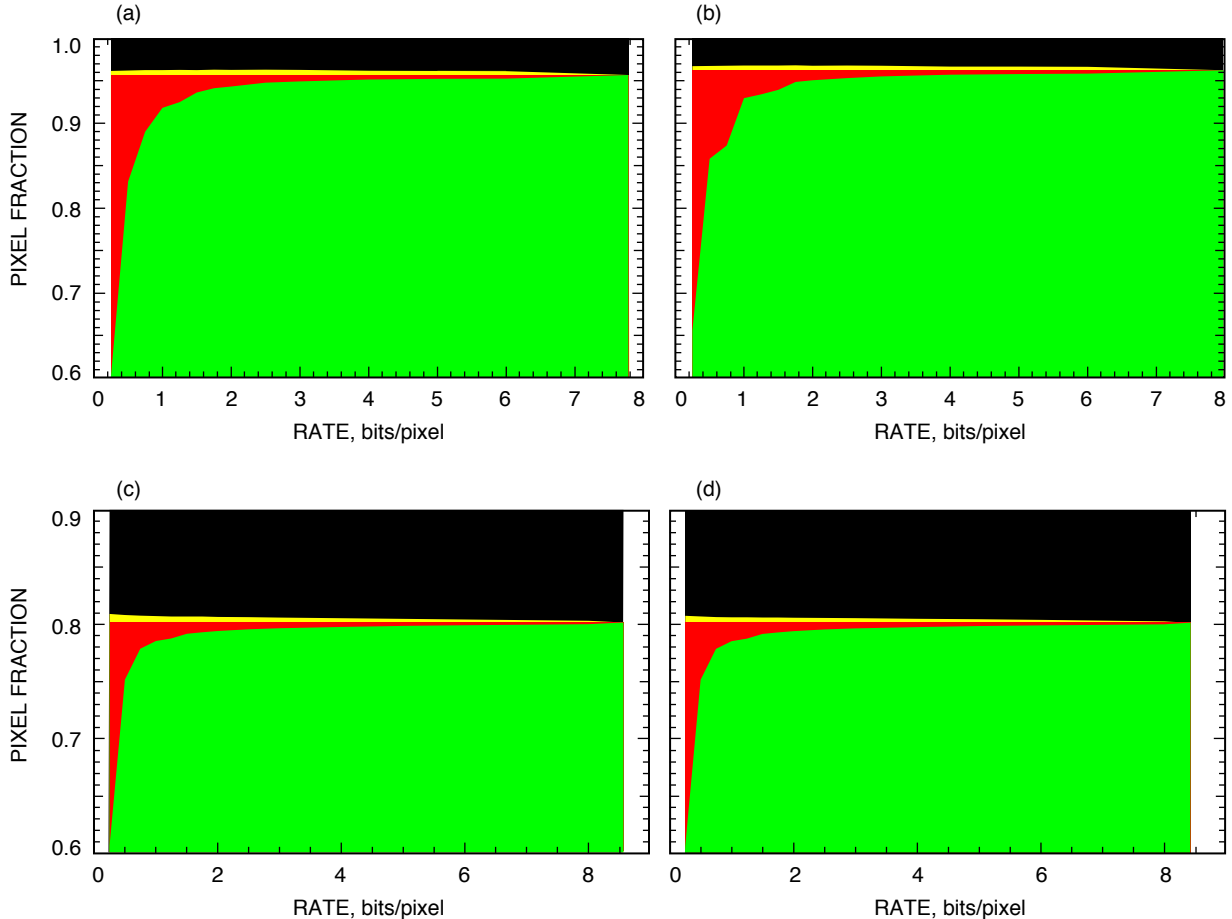


Fig. 15. For each test point, a curve shows the discrepancy between surveyed and stereo-estimated range (averaged over the corresponding 9-pixel cluster) as a function of compressed bit rate in (a) HazCam scene A, (b) HazCam scene B, (c) NFOV-Cam scene A, and (d) NFOV-Cam scene B.



**Fig. 16. Fraction of linearized and rectified image pixels for which range estimates are produced under lossy and lossless compression (green), lost (red), found (yellow), no range estimates (black) as a function of bit rate for (a) high-interest pixels in HazCam scene A, (b) high-interest pixels in HazCam scene b, (c) NFOV-Cam scene A, and (d) NFOV-Cam scene B.**

Figure 16 suggests that the fraction of found pixels is small for NFOV-Cam scenes and high-interest HazCam pixels, and the fraction of lost pixels becomes significant at low bit rates.

We say that a pixel is *available* if the stereo ranging process produces a range estimate for that pixel when the original image pair is used. The fraction of available pixels that becomes lost provides a rough indication of the quality of the stereo ranging data product relative to that obtained under lossless compression, but it tends to understate the quality degradation at very low bit rates, where a significant fraction of the range estimates produced may be poor. Thus, we'd like to also count the fraction of available pixels whose range estimates have error exceeding some threshold. Unfortunately, measuring true range values for more than a small number of points in a scene is impractical. As a crude approximation, for each available pixel we compute  $\hat{\sigma}$ , an estimate of range uncertainty  $\sigma$  based on the range results for the original image pair. For a given lossy-compressed stereo pair, we say that a pixel is *unreliable* if the range estimate differs from the lossless estimate by more than  $\hat{\sigma}$ . This method of classifying pixels as unreliable is convenient, but has the drawback that we have no practical way to compare it against ground truth.

Figure 17 shows the fraction of available NFOV-Cam and high-interest HazCam pixels that becomes lost or unreliable. At low bit rates, evidently lost pixels are more of a problem for HazCam and unreliable pixels are more of a problem for NFOV-Cam scenes. Comparison of Figs. 17 and 15 suggests that the

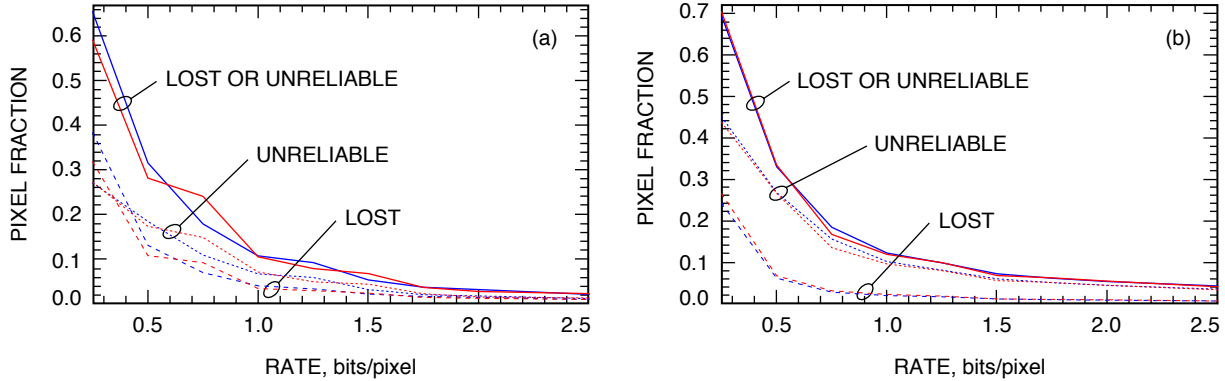


Fig. 17. Fraction of available pixels that are lost or unreliable for (a) high-interest pixels HazCam scenes A and B and (b) NFOV-Cam scenes A and B. Curves for scene A are shown in blue, scene B in red.

fraction of unreliable pixels may overstate somewhat the loss of range accuracy at low bit rates. The fraction of available pixels that become lost or unreliable provides a rough indication of stereo quality. This metric is similar to the one proposed by Nagura [8–10]. For HazCam scenes, it is reasonable to limit our consideration to high-interest pixels.

## V. Modifying ICER

In this section, we consider how to modify ICER to improve stereo ranging data quality obtained under lossy compression even *without* jointly compressing the two images in a stereo pair. The idea is to change which portion of the discrete wavelet transform (DWT) data is encoded by ICER at a given compressed bit rate.

### A. Subband Bit-Plane Prioritization

ICER output consists of compressed bit planes of wavelet-transformed image subbands [5]. Each successive compressed subband bit plane improves reconstructed image quality, so to optimize the trade-off between compressed bit rate and image quality, at each stage in compression we’d like to select for encoding the subband bit plane that gives the largest improvement in reconstructed image quality per compressed bit. However, we’d like to avoid a computationally expensive image-dependent optimization to make this selection, and so ICER relies on a simple heuristic rule to assign a priority value to each subband bit plane. The prioritization rule is intended to assign higher priority values to bit planes that provide larger reductions in MSE distortion per compressed bit.

As an example, Fig. 18 plots MSE distortion as a function of compressed bit rate for HazCam image  $\Psi$  and NavCam image  $\Phi$ . For each compressed subband bit plane, we measure the resulting reduction in MSE distortion,  $\Delta\text{MSE}$ , and the cost in compressed bits,  $\Delta b$ . The ratio  $\Delta\text{MSE}/\Delta b$  gives the average reduction in MSE distortion per compressed bit for the subband bit plane. Figure 19 plots  $\Delta\text{MSE}/\Delta b$  as a function of cumulative bit rate for these two images, and Fig. 20 plots this ratio as a function of priority value for each subband of the images.

Figure 19 indicates that ICER’s simple prioritization rule is reasonably effective at optimizing MSE distortion as a function of bit rate. The graph is nearly monotonic, and subband bit planes having the same priority value tend to give about the same reduction in MSE distortion per transmitted bit. Places where re-ordering bit planes would improve rate-distortion performance generally correspond to small segments of compressed data, suggesting that changing the bit-plane priority rule used by ICER would provide only a very modest benefit when MSE distortion is the quality metric of interest.<sup>7</sup>

<sup>7</sup>The number of compressed bits used to encode a subband bit plane depends on the order in which the bit planes are encoded, but in practice this effect is negligible.

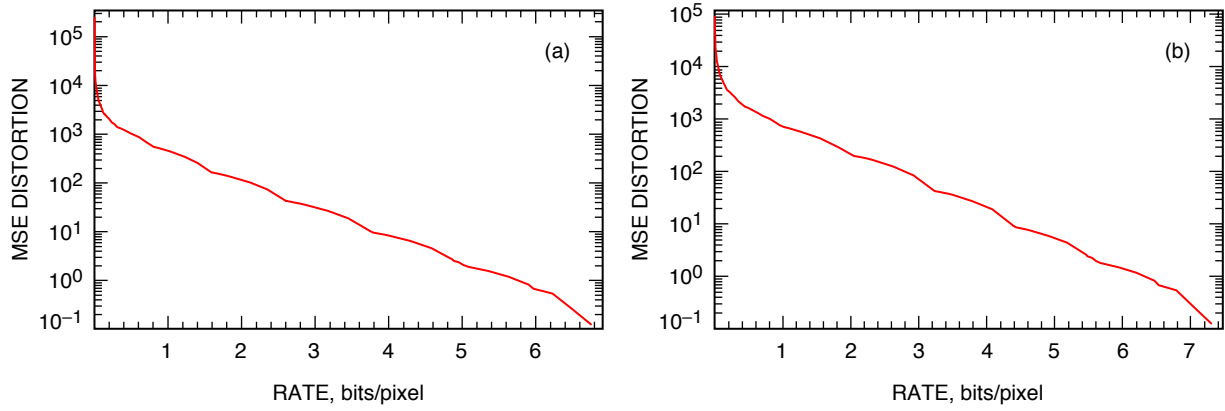


Fig. 18. Rate-distortion curves for (a) HazCam image  $\Psi$  and (b) NavCam image  $\Phi$ .

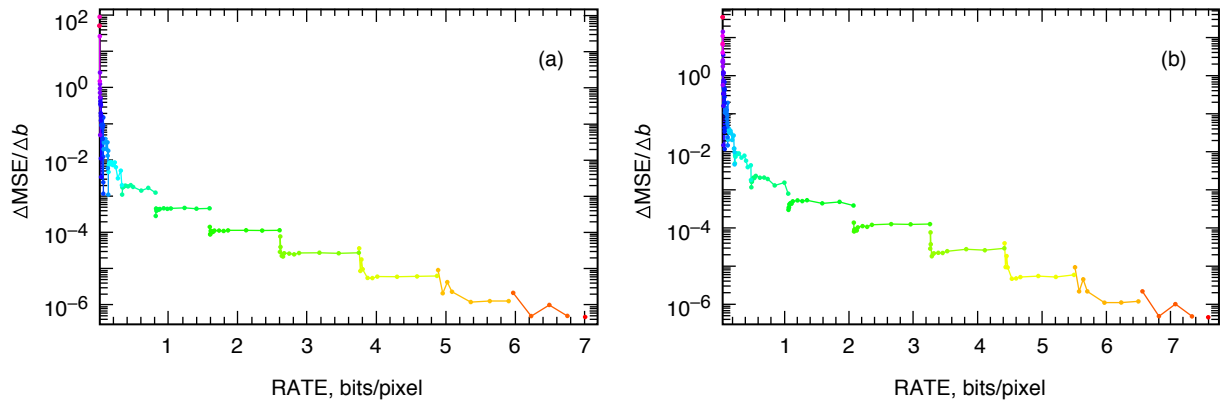


Fig. 19.  $\Delta\text{MSE}/\Delta b$  versus bit rate for (a) HazCam image  $\Psi$  and (b) NavCam image  $\Phi$ . Different colors indicate different priority values, as assigned by ICER to the subband bit plane.

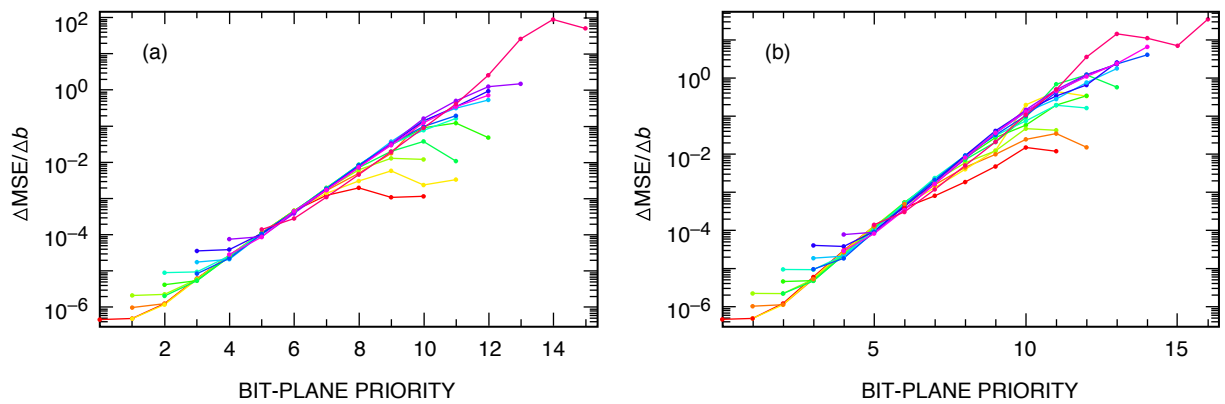


Fig. 20.  $\Delta\text{MSE}/\Delta b$  versus bit-plane priority for each subband bit plane of (a) HazCam image  $\Psi$  and (b) NavCam image  $\Phi$ . Each subband is shown in a different color.

Figure 20 further demonstrates that subband bit planes with the same priority value tend to have approximately the same value of  $\Delta\text{MSE}/\Delta b$ . Exceptions arise for some of the most significant bit planes and the last bit plane of each subband. But the most significant bit planes tend to be very economical to compress anyway, and ICER’s rule for breaking ties among bit planes with equal priority fortunately causes the last bit plane of each subband to be correctly prioritized.

## B. Alternative Distortion Metric

As described in Section I.A, before the stereo ranging process produces a disparity map, images undergo linearization, rectification, and box filtering. By considering the impact of these operations, we can (1) determine a more relevant quality metric than MSE distortion and (2) determine an alternative bit-plane prioritization rule that does a better job of optimizing this metric and, in turn, produces better ranging data products. Ideally, we’d like to directly consider the impact on ranging results (e.g., the fraction of pixels that is either lost or unreliable) of each successive subband bit plane, but this approach did not seem practical.

ICER compresses bit planes of subbands produced by a DWT, and so it is informative to view the impact of linearization, rectification, and box filtering in the DWT domain. That is, we would like to compare the DWT domain representation of the original image with the DWT domain representation of the image following linearization, rectification, box filtering, and then inversion of the linearization/rectification process. Figures 21 and 22 provide such a comparison for HazCam image  $\Psi$  and NavCam image  $\Phi$ . Images are cropped after linearization and rectification, resulting in the loss of some information near the edges of images. This effect is conspicuous in the HazCam image but not in the NavCam image. More significantly, the box-filtering operation smoothes the linearized/rectified image, and we see from the figures that the high spatial frequency subbands have much lower energy. This suggests that, when the images are to be used for stereo ranging, lossy compression effectiveness might be improved by altering the prioritization scheme to reduce the priority of the corresponding high spatial frequency subbands, since much of the information in these subbands is effectively filtered out.

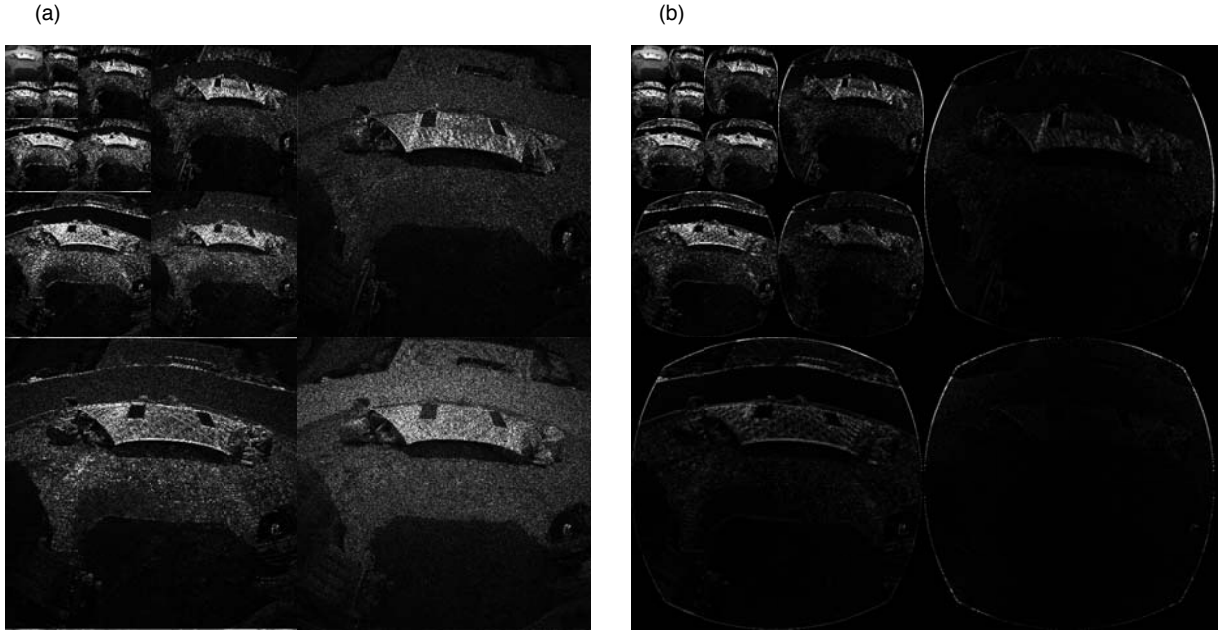
To take into account the processing performed on the images, it seems appropriate to measure distortion between an original and reconstructed image *after* the linearization/rectification and box-filtering operations have been performed on both images. Let  $I$  and  $\hat{I}$  denote original and reconstructed images, and let  $LR$  and  $B$  denote the linearization/rectification and box-filtering operations, respectively. Our stereo-motivated distortion metric is

$$D_s(I, \hat{I}) = \text{MSE} \left[ B(LR(I)), B(LR(\hat{I})) \right]$$

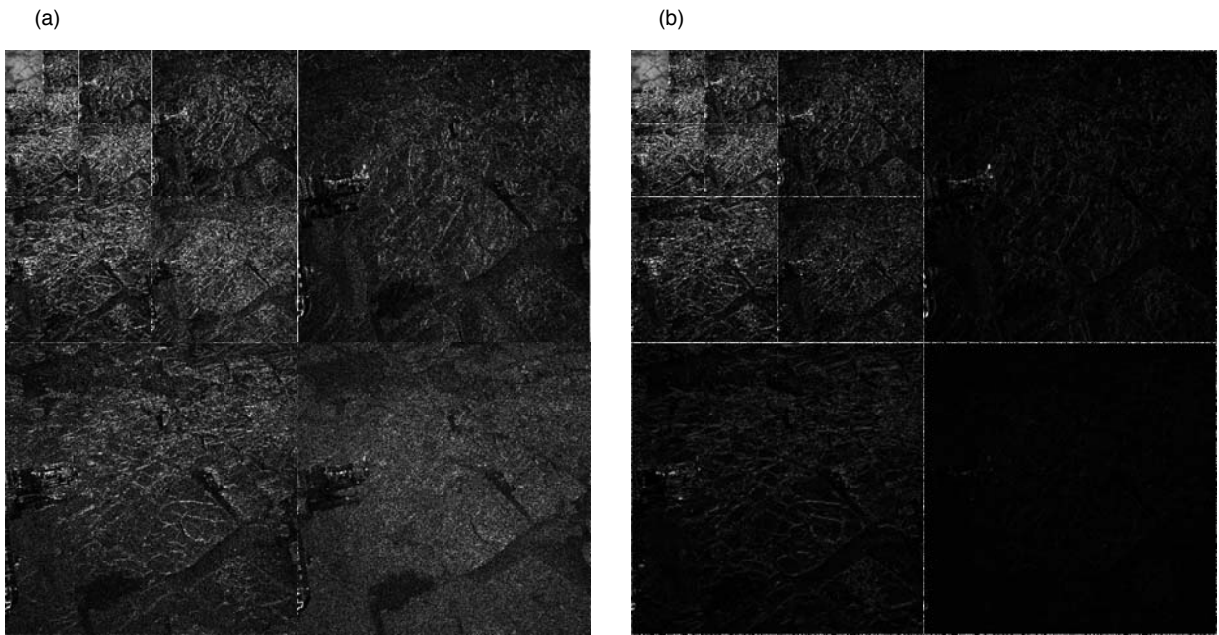
For each successive subband bit plane, we measure the reduction in this distortion,  $\Delta D_s$ , and the cost in compressed bits,  $\Delta b$ . For HazCam image  $\Psi$  and NavCam image  $\Phi$ , Fig. 23 shows the ratio  $\Delta D_s/\Delta b$  as a function of bit rate and Fig. 24 shows  $\Delta D_s/\Delta b$  as a function of bit-plane priority value for each subband. We see from Fig. 23 that subband bit planes assigned the same priority by ICER may have very different values of  $\Delta D_s/\Delta b$ , which suggests that an alternative priority assignment scheme might lead to improved stereo ranging performance. Figure 24 further demonstrates that to improve rate-distortion performance (with respect to distortion metric  $D_s$  on these sample images) the bit-plane prioritization scheme should be adjusted to give higher priority to certain subbands.

## C. Priority-Adjusted ICER

Based on the results of Fig. 24, we define a priority adjustment value for each subband. Table 6 gives the priority adjustments using the subband naming convention described in [5]. To make use of Table 6, during compression each subband bit-plane priority is increased by the adjustment value listed in



**Fig. 21.** DWT-domain representations of HazCam  $\Psi$ : (a) DWT of original image and (b) DWT after linearization/rectification, box filtering, and then inversion of the linearization/rectification process. Subbands show DWT coefficient magnitudes, scaled by a factor of 8 relative to the lowest spatial-frequency subband.



**Fig. 22.** DWT-domain representations of HazCam  $\Phi$ : (a) DWT of original image and (b) DWT after linearization/rectification, box filtering, and then inversion of the linearization/rectification process. Subbands show DWT coefficient magnitudes, scaled by a factor of 8 relative to the lowest spatial-frequency subband.

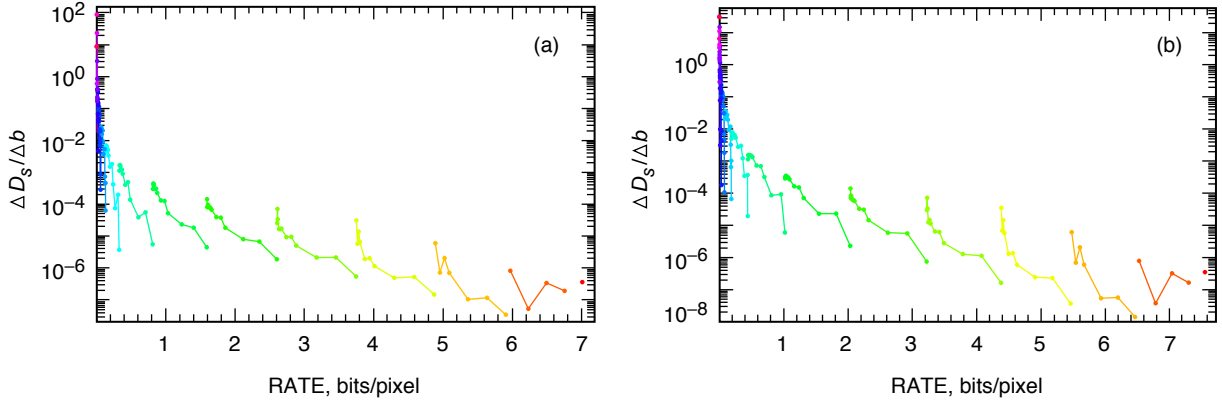


Fig. 23.  $\Delta D_s/\Delta b$  versus bit rate for (a) MER HazCam image  $\Psi$  and (b) MER NavCam image  $\Phi$ . Different colors indicate different priority values, as assigned by ICER to the subband bit plane.

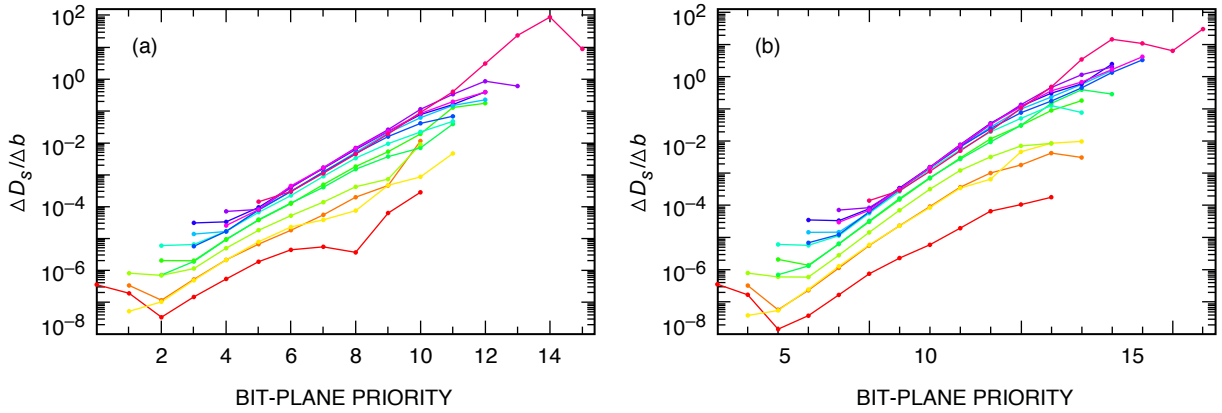


Fig. 24.  $\Delta D_s/\Delta b$  versus bit-plane priority for each subband bit plane of (a) MER HazCam image  $\Psi$  and (b) MER NavCam image  $\Phi$ . Each subband is shown in a different color.

Table 6. Subband priority adjustments. See [5] for a description of the labels used to identify subbands.

Subband	Priority adjustment
Level 1 HH	0
Level 1 LH, HL	1.35
Level 2 HH	2
Level 2 LH, HL	2.5
Level 3 HH	3
Others	3.1

the table. For HazCam image  $\Psi$  and NavCam image  $\Phi$ , Fig. 25 shows that the adjusted priority values provide a much better indicator of reduction in the distortion  $D_s$  per compressed bit than do the original priority values as illustrated in Fig. 24.

For HazCam image  $\Psi$  and NavCam image  $\Phi$ , Fig. 26 shows the distortion  $D_s$  as a function of compressed bit rate under regular and priority-adjusted versions of ICER, along with a bound on  $D_s$  obtained by sorting the values of  $\Delta D_s/\Delta b$  (ignoring the fact that bit planes within a subband must be transmitted in sequence). The graph shows that the priority-adjusted version of ICER provides distortion quite close to the bound, and thus further refinements of the subband prioritization rule would make little improvement on rate-distortion performance for these images under the distortion measure  $D_s$ . Figure 27 shows a similar improvement in distortion  $D_s$  as a function of compressed bit rate for HazCam scene  $\Omega$  and NavCam scene  $\Theta$ .

Finally, we evaluate the effect of the priority adjustment on stereo ranging quality, as measured by the combined fraction of lost and unreliable pixels, for HazCam scene  $\Omega$  and NavCam scene  $\Theta$ . Figure 28 illustrates that, under this stereo quality metric, priority-adjusted ICER performs at least as well as regular ICER for these scenes, and in some instances provides a significant improvement. For example, on NavCam scene  $\Theta$ , under regular ICER at 1 bit/pixel, 16 percent of the available pixels are lost or unreliable. Under priority-adjusted ICER, this level of quality is achieved at 0.7 bits/pixel—this corresponds to a 30 percent reduction in bit rate required to achieve a given quality level.

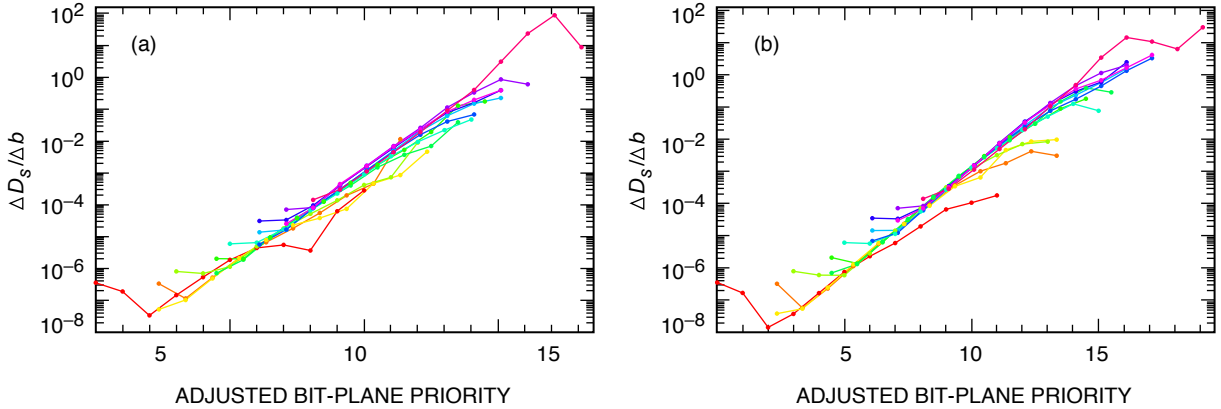


Fig. 25.  $\Delta D_s/\Delta b$  versus *adjusted* bit-plane priority for each subband bit plane of (a) HazCam image  $\Psi$  and (b) NavCam image  $\Phi$ . Each subband is shown in a different color.

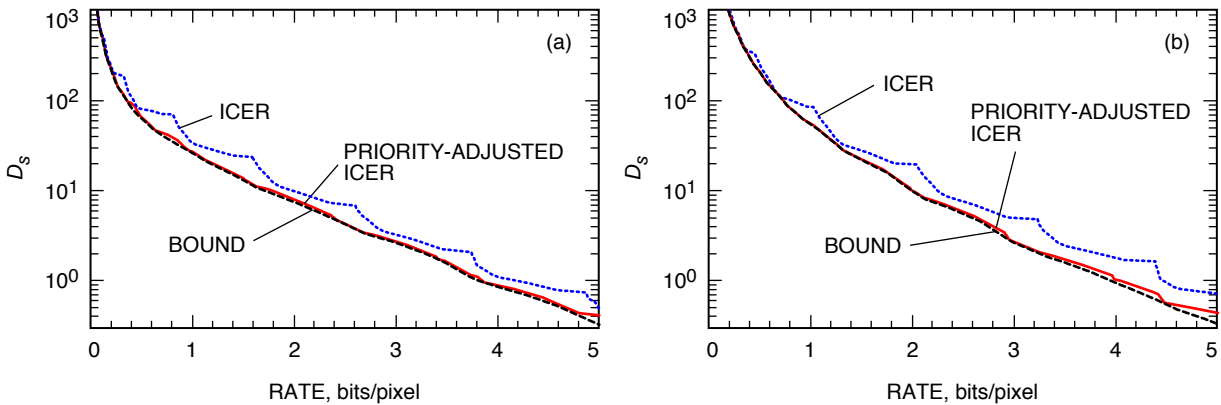


Fig. 26.  $D_s$  versus compressed bit rate for (a) HazCam image  $\Psi$  and (b) NavCam image  $\Phi$ .



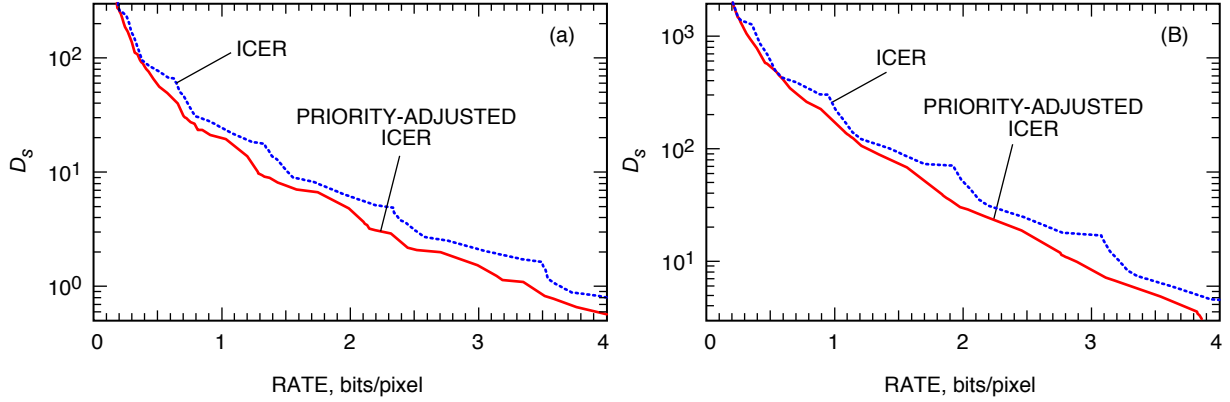


Fig. 27.  $D_s$  versus compressed bit rate for (a) high-interest pixels in the left image of HazCam image  $\Omega$  and (b) the left image of NavCam image  $\Theta$ .

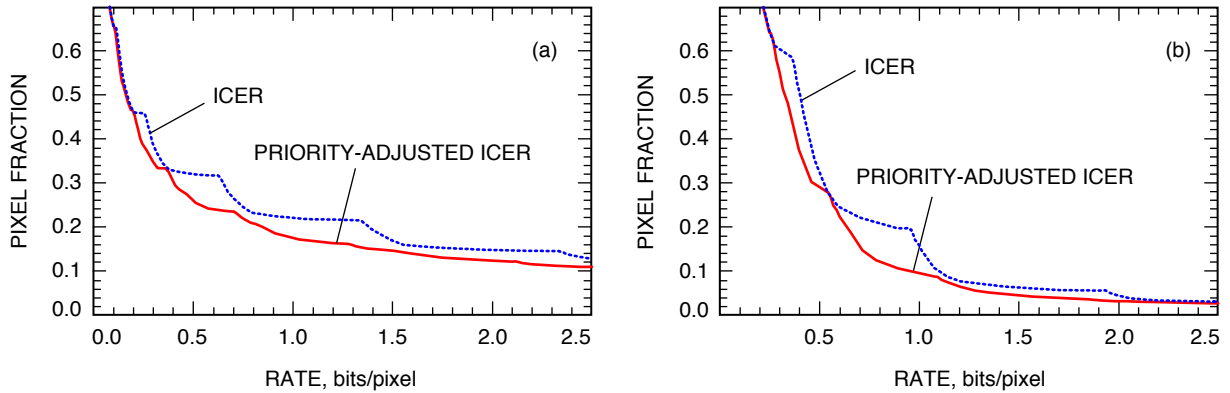


Fig. 28. Fraction of available pixels that become lost or unreliable for (a) high-interest pixels in HazCam scene  $\Omega$  and (b) NavCam scene  $\Theta$ .

## VI. Conclusion

One of the purposes of our investigations was to determine the effects of lossy compression on stereo ranging. These effects can be measured directly, using ground truth data, or indirectly, by considering first the uncertainty in stereo ranging using the original images, then adding the uncertainty attributable to lossy compression. Our direct measurements of the effects of lossy compression on stereo ranging with respect to ground truth suggest that at bit rates above about 2 bits/pixel, changes in ranging accuracy for the test points are negligible compared to the error inherent in the stereo ranging process. These measurements also essentially agreed with known models of stereo-derived ranging uncertainty for the original images, although they suggest that the stereo-matching algorithm might achieve disparity accuracy slightly better than the assumed 1/3 pixel. In any case, it is clear that some points can deviate significantly from the model in certain circumstances that make stereo matching challenging (saturated pixels, object edges, etc.).

We also took the indirect route to determining the effects of lossy compression on stereo ranging. In this case, our evaluation considers both lost pixels (range estimate no longer found) and unreliable pixels (range estimate differs by more than the inherent stereo-matching uncertainty). Here fairly extensive quantitative results were obtained on how the fractions of lost and unreliable pixels increase as the compressed bit rate decreases. It is difficult to make non-obvious characterizations of the behavior, but

we hope that the quantitative results will give rover operations planners additional data upon which to judge the degree of lossy compression to use.

The box-filtering step performed as part of the stereo ranging process effectively diminishes the impact of high spatial frequency data on stereo ranging results. With this in mind, we defined an image distortion metric,  $D_s$ , that takes into account the image-processing steps performed as part of the stereo ranging process and provides a better indication of stereo ranging quality than MSE distortion. The distortion metric is somewhat cumbersome to compute (it requires linearizing images and then box filtering the result), but it serves as a useful tool to analyze ICER's bit-plane prioritization scheme.

By evaluating the empirical impact of subband bit planes on the distortion metric  $D_s$ , we developed a modification to ICER's bit-plane prioritization rule that improves stereo ranging quality as a function of compressed bit rate. The modified prioritization rule could easily be implemented as an option in MSL compression flight software; it is simple to implement and does not increase compression complexity. In our tests, the modified prioritization rule frequently allowed a given fraction of lost or unreliable pixels to be reached at a significantly lower bit rate than with the original ICER (sometimes 30 percent or more), although there was wide variation. The modified prioritization rule never made the results worse.

We note that our modification to ICER does not exploit correlations between the left and right images in a stereo pair. In most cases, the gains achieved by modifying ICER to exploit such correlation would be cumulative with the improvement from our priority-adjusted ICER.

## Acknowledgment

The authors would like to thank Bob Deen for providing software needed to duplicate the MER stereo ranging process and for several helpful discussions about its operation and inner workings.

## References

- [1] S. D. Blostein and T. S. Huang, "Error Analysis on Stereo Determination of 3-D Point Positions," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 9, no. 6, pp. 752–765, 1987.
- [2] R. G. Deen and J. J. Lorre, "Seeing in Three Dimensions: Correlation and Triangulation of Mars Exploration Rover Imagery," *2005 IEEE International Conference on Systems, Man and Cybernetics*, vol. 1, Waikoloa, Hawaii, pp. 911–916, October 10–12, 2005.
- [3] D. B. Gennery, "Generalized Camera Calibration Including Fish-Eye Lenses," *International Journal of Computer Vision*, vol. 68, no. 3, pp. 239–266, 2006.
- [4] S. B. Goldberg, M. W. Maimone, and L. Matthies, "Stereo Vision and Rover Navigation Software for Planetary Exploration," *2002 IEEE Aerospace Conference Proceedings*, vol. 5, Big Sky, Montana, pp. 5-2025–5-2036, March 2002.

- [5] A. Kiely and M. Klimesh, "The ICER Progressive Wavelet Image Compressor," *The Interplanetary Network Progress Report 42-155, July–September 2003*, Jet Propulsion Laboratory, Pasadena, California, pp. 1–46, November 15, 2003. [http://ipnpr.jpl.nasa.gov/progress\\_report/42-155/155J.pdf](http://ipnpr.jpl.nasa.gov/progress_report/42-155/155J.pdf)
- [6] A. Kiely and M. Klimesh, "Preliminary Image Compression Results from the Mars Exploration Rovers," *The Interplanetary Network Progress Report*, vol. 42-156, Jet Propulsion Laboratory, Pasadena, California, pp. 1–8, February 15, 2004. [http://ipnpr.jpl.nasa.gov/progress\\_report/42-156/156I.pdf](http://ipnpr.jpl.nasa.gov/progress_report/42-156/156I.pdf)
- [7] J. N. Maki, J. F. Bell III, K. E. Herkenhoff, S. W. Squyres, A. Kiely, M. Klimesh, M. Schwochert, T. Litwin, R. Willson, A. Johnson, M. Maimone, E. Baumgartner, A. Collins, M. Wadsworth, S. T. Elliot, A. Dingizian, D. Brown, E. C. Hagerott, L. Scherr, R. Deen, D. Alexander, and J. Lorre, "The Mars Exploration Rover Engineering Cameras," *Journal of Geophysical Research—Planets*, vol. 108, no. E12, 8071, December 2003.
- [8] R. Nagura, "Multi-Stereo Imaging System using a New Data Compression Method," *IEEE 2006 International Geoscience and Remote Sensing Symposium (IGARSS '06)*, Denver, Colorado, July 31–August 4, 2006.
- [9] R. Nagura, "Optimum Data Transmission and Imaging Method for High Resolution Imaging from Earth Observation Satellite," *2003 IEEE Workshop on Advances in Techniques for Analysis of Remotely Sensed Data*, pp. 163–172, October 27–28, 2003.
- [10] R. Nagura, "Accurate 3D Information Extraction from Large-Scale Data Compressed Image and the Study of the Optimum Stereo Imaging Method," *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. XXXIV, part 5/W10, ISPRS Workshop on Visualization and Animation of Reality-based 3D Models, Tarasp, Switzerland, pp. 1–9, February 24–28, 2003.
- [11] D. S. Taubman and M. W. Marcellin, *JPEG2000: Image Compression Fundamentals, Standards and Practice*, Boston, Massachusetts: Kluwer Academic Publishers, 2002.
- [12] M. J. Weinberger, G. Seroussi, and G. Sapiro, "The LOCO-I Lossless Image Compression Algorithm: Principles and Standardization into JPEG-LS," *IEEE Transactions on Image Processing*, vol. 9, no. 8, pp. 1309–1324, August 2000.
- [13] Y. Yakimovski and R. T. Cunningham, "A System for Extracting Three-Dimensional Measurements from a Stereo Pair of TV Cameras," *Computer Graphics and Image Processing*, vol. 7, pp. 195–210, 1978.