Integer Cosine Transform Compression for Galileo at Jupiter: A Preliminary Look

L. Ekroot, S. Dolinar, and K.-M. Cheung
Communications Systems Research Section

The Galileo low-gain antenna mission has a severely rate-constrained channel over which we wish to send large amounts of information. Because of this link pressure, compression techniques for image and other data are being selected. The compression technique that will be used for images is the integer cosine transform (ICT). This article investigates the compression performance of Galileo's ICT algorithm as applied to Galileo images taken during the early portion of the mission and to images that simulate those expected from the encounter at Jupiter.

I. Introduction

The Galileo low-gain antenna (LGA) mission will rely heavily on source coding, i.e., data compression, to maximize the return over the severely rate-constrained channel. A major portion of Galileo's information return from Jupiter will be images from the solid-state imaging (SSI) camera. Image frames contain either 800 × 800 or 400 × 400 8-bit pixels, and would require 5.12 Mbits or 1.28 Mbits, respectively, to transmit uncompressed. Image compression gains of 10 dB and higher are obtainable through known coding techniques such as the discrete cosine transform (DCT), which introduces hardly noticeable losses in data fidelity. However, Galileo cannot directly utilize off-the-shelf compression algorithms due to its limited onboard computing capability. Instead, Galileo will use a computationally tractable algorithm called the integer cosine transform (ICT). The ICT operates and performs comparably to the DCT. This article investigates the compression performance of Galileo's ICT algorithm, as applied to actual Galileo images and to simulated images constructed to mimic the conditions that the Galileo camera is expected to encounter at Jupiter.

The amount of compression achieved by the ICT is adjustable by selecting the degree to which the transform coefficients are quantized. Coarser quantization generally produces higher compression and larger errors in the reconstructed image. Thus, the study of ICT compression performance is in fact an analysis of the possible tradeoffs between compression of a source image and the resulting distortions in the reconstructed image. It is important to understand the extent to which the compression is predictable by using anticipated image characteristics, and how compression versus distortion varies with the selectable quantization step size. The overall statistical variability of the compression performance is an important

factor that influences link usage planning. This article is a preliminary attempt to examine these issues.

The ICT is described in [1] and [2], and the specific application of the ICT to the Galileo LGA mission is discussed in [3]. For a given quantization step size, quantization matrix or template, Huffman table, and, of course, image, the ICT produces a compressed image. The compression ratio indicates the ratio of the compressed image size to the original image size. Distortion indicates the difference between the reconstructed image and the original one. Neither the compression ratio nor the distortion is input for the ICT. This article describes the behavior of the compression ratio and distortion as a function of the quantization step size, q, and any prior information about the image.

An additional level of control may be achievable by selecting one of a few specially constructed quantization templates and using a corresponding Huffman table. For this article, the algorithm uses only a template for uniform quantization and a Huffman table that has been trained for the template on space images [3]. The tailoring and use of different quantization matrices is the subject of ongoing work, and will not be discussed here.

II. The Images in the Data Set

The ICT has been tested on various images. Ideally, the test images would be ones of Jupiter taken by Galileo's SSI camera. Instead, the available images are ones taken by Galileo of Venus, the Earth, the Moon, and black sky, and ones taken by Voyager of Jupiter and a few of Jupiter's satellites. There are four image sets provided for this study by JPL's Multimission Image Processing Laboratory (MIPL).

(1) The Earth 1 images are Galileo images from volumes 2 through 6 of the Earth 1 CD ROMs. They consist of 970 raw_cal, 1798 Earth, 81 Venus, and 773 Moon images for a total of 3622 images. A randomly selected one-third of the Earth 1 images has been used for this article to identify trends and performance predictors. The other two-thirds of the data has been reserved for future determination of the effectiveness of the predictors and the consistency of the trends.

(2) The 13 selected Galileo SSI images consist of 12 recent Galileo images and 1 of the Earth 1 images selected by MIPL to illustrate a typical range of image types. These images are designated by names beginning with rq534.

(3) The 19 simulated Galileo images were obtained from Voyager images of the Jovian system by removing Voyager camera characteristics and introducing the Galileo SSI camera signature. These images are designated by names r.1, r.2, ..., r.19.

(4) The 16 radiation-noise-added simulated Galileo images are 7 of the simulated images and 9 of the selected images with 1 of 4 available noise frames (A, B, C, and D) added by MIPL to simulate the radiation effects at Jupiter. These images are designated by names beginning with rq538.

Some examples of the selected Galileo images are shown in Fig. 1. The smaller data sets are useful for illustrating different behaviors but do not provide sufficient data for analysis. The Earth 1 images provide a larger number of images for drawing statistical inferences.

Images were compressed using different values of q, and reconstructed. Quantization step sizes 2 through 18, 20, and 24 were tested on all images. The compression ratio, root mean squared error (rmse), error variance, and error mean were noted along with image statistics such as the mean, minimum and maximum pixel values, the sample zeroth-order entropy, and the sample difference entropy. From this information and some a priori information about the image, patterns in the compression and distortion performance have been sought.

III. Performance as a Function of Quantization

This section shows the relationships between compression ratio and image distortion for the ICT as the quantization step size is varied.

A. Compression Behavior

The compression ratio achieved by the ICT depends on the image. This section shows the gross behavior with only the small representative image sets. Scatter plots illustrate the compression behavior as a function of the quantization step size, the spread of compression ratios for different images, the limiting compression for very large q, and the effect of the simulated radiation noise on the compression.

For the Huffman codes used in this work, there is a limit of 102.4 on the compression ratio achievable by the 8×8 ICT. This is because an all-zero quantized block in the transform domain can be coded in 5 bits, and an

uncompressed block is 8×64 bits. In order for a quantized block to be all zero, either the original block had to be nearly black or the quantization step size had to be very large. In either case, the reconstructed block would be entirely black. Among the images tested, the nearly black ones approached this limit with the quantization step sizes tested.

Figure 2 has three scatter plots showing the compression ratio versus quantization step size for the images in each of the three small data sets. The plots in Fig. 2 show that the compression trend is typically monotonic, increasing with q, but that it starts to level off for higher q. In most instances, as q increases, the spread on the log scale of the compression ratio decreases; this is consistent with the upper limit on compression ratio. The unlabeled scatter plots also show that without taking into account information about the specific image, it is difficult to predict the compression.

To illustrate the effects on compression of adding the simulated radiation noise to images, Fig. 3 has one plot for each of the four noise frames used. The plots show how the noise changed the compression ratio for the individual images. Solid symbols are used for images in the sets of selected Galileo SSI images or simulated Galileo images, and hollow symbols with the same shape are used for the corresponding images with radiation noise added. These plots seem to indicate that the noise frames dominate the compression performance; this is more apparent for frames A and B and less so for D.

B. Distortion

The rmse is a measure of the pixel-by-pixel difference between the original image and the reconstructed image. It is not the best measure of the distortion of an image, but it is easily quantifiable and commonly used as a guideline. A better measure of the distortion is how it affects the conclusions and analysis of the images sent back from Jupiter, but this is hard to quantify. The SSI team, consisting of scientists who will be using Galileo's images, has been meeting with both JPL and the Ames Research Center to try to determine what types of distortion are of more concern than others. The appropriate distortion measure for this analysis is not the subject of this article, and for now the rmse will serve as a rough guideline.

For a given image, the rmse typically increases as the quantization step size increases. The three scatter plots in Fig. 4 show rmse versus quantization step size for the three smaller image sets. Although there is a monotonic trend, there are examples of images for which the rmse is unexpectedly high for a particular q or q's. These high rmse situations may be avoided most of the time by careful selection of the quantization step size.

C. Selection of Quantization Step Size

For a particular image, there may be a quantization step size that is inferior, meaning that for that image there is another quantization step size that gives both better compression and better rmse. An inferior q causes a non-monotonic spike on the compression versus distortion curve. The plot in Fig. 5 shows the distortion versus compression ratio for the selected Galileo SSI images in Fig. 1, namely rq534.litn, rq534.ausvn, rq534.lun, rq534.pct, rq534.ear9n, and rq534.gas. The solid lines show the curves traced out by the rmse and compression ratios at all of the tested q's; the dashed lines show how the curves would be smoother and more predictable if the inferior q's could be predicted and avoided. Images that are predominately black sky, such as the image of Gaspra, rq534.gas, exhibit the most dramatic rmse fluctuations with q.

In order to get a reasonable profile for quantization step sizes that are likely to be inferior, the larger image set is necessary. The number of images for which a particular q was inferior was tallied over the randomly selected third of the Earth 1 images. Figure 6 is a histogram showing the frequency that different q's were inferior for the raw_cal or black sky images, and for the non-raw_cal images that are of the Earth, Moon, and Venus. The histogram shows that some q's are inferior more often than others, but sometimes only for certain subsets of the images. For instance, $q = 6$ is inferior for a quarter of the raw_cal images, but not often for the interesting ones. For $q = 13$, we see the opposite trend. Further study is important and should include a measure of how bad a q is for a particular image (as opposed to simply counting how often q occurs). Also, if the images will be classified into types that give improved prediction of the compression, then the study of which q's are inferior should be redone for each subset.

IV. Performance as a Function of Image Characteristics

Using the randomly selected one-third of the Earth 1 images, correlations were done between the camera settings and the compression ratio, the compressed image size, and the rmse in order to direct the study towards likely parameters for predictions, or at least away from

parameters that appear to be uncorrelated. The camera settings were not continuously valued, and so were not ideal for a correlation analysis. A portion of the correlation matrix of interest is shown in Fig. 7.

The filter width has a small amount of correlation with the performance. The middle wavelength of the camera filters seemed independent of the compression and the rmse. By dividing the selected data set into subsets defined by the different camera settings, statistics for each subset could be found. Scatter plots of compression versus rmse show the trends for different subsets. However, an analysis of variance test with a level of significance of $\alpha = 0.05$ indicated that the subsets probably come from distributions with the same mean.

The black sky images may need to be treated separately because compression performance effects that are second order for more interesting images are apparently first order for these images. For instance, the raw_cal images that are mostly black sky tend to cluster according to the gain state on the camera (not shown in this article). The effect is less noticeable for images with less black sky.

V. Overall Statistical Variability

With so many images in the Earth 1 image set, the compression and rmse distortion behavior can be displayed in histograms where the range of values is divided into bins and in cumulative distribution type plots. Figures 8 and 9 show the histograms and cumulative distributions of both compression ratio and rmse for the selected third of the images using $q = 8$. On the cumulative plots there are three curves: one is for the Earth, Moon, and Venus images; one is for the raw_cal images; and one is for all of the images. There are separate histograms corresponding to each curve on the cumulative plots. Histograms of this sort for different quantization step sizes can assist Galileo mission planners in deciding how many pictures to attempt to send back, and at what risk this can be done.

VI. Conclusions

This preliminary study has generally validated the assumptions regarding high image compression gains that the Galileo LGA mission is counting on. But it has also flagged several areas of concern that require further study.

The scatter plots of Figs. 2–4 and the histograms of Fig. 8 show that compression ratios of 10:1 to 40:1 are not unreasonable to expect at moderately small values of rmse. On the other hand, there is a statistical variability of the achievable compression ratio on the order of 5 to 10 dB from image to image, and slightly less variability in the rmse. The variability of compressibility can greatly complicate the job of planning how many images can be sent back and how large buffers have to be in the onboard computer.

The study of correlations between the known camera settings and the resulting compression performance suggested no significant correlation that could be used to appreciably improve the predictability of performance.

There needs to be more analysis of what can be calculated or predicted about each image to aid downlink planning and, perhaps, compression control. For instance, better predictors might be obtained from anticipatable or measurable image characteristics, such as the percentage of black sky expected to fill the image frame. Alternatively, an onboard calculation could be used to adjust the quantization step size over small portions of the image in order to target a particular compression ratio. An example of such a calculation of local image statistics is an estimate of the entropy using the buffered portion of the image awaiting compression. Such ideas are under examination.

Further improvement in the reconstructed images can be achieved by developing more sophisticated reconstruction techniques to exploit continuities across block boundaries and other spatial correlations in the image. There is ongoing work to develop special purpose quantization matrices for improved image quality at good compression ratios.

References

[1] W. Cham, "Development of Integer Cosine Transform by the Principle of Dyadic Symmetry," *IEE Proceedings*, vol. 136, no. 4, pp. 276–282, August 1989.

[2] K.-M. Cheung, F. Pollara, and M. Shahshahani, "Integer Cosine Transform for Image Compression," *The Telecommunications and Data Acquisition Progress Report 42-105, vol. January–March 1991*, Jet Propulsion Laboratory, Pasadena, California, pp. 45–53, May 15, 1991.

[3] K.-M. Cheung and K. Tong, "Proposed Data Compression Schemes for the Galileo S-Band Contingency Mission," *1993 Space and Earth Science Data Compression Workshop Proceedings*, Snowbird, Utah, pp. 99–109, April 2, 1993.

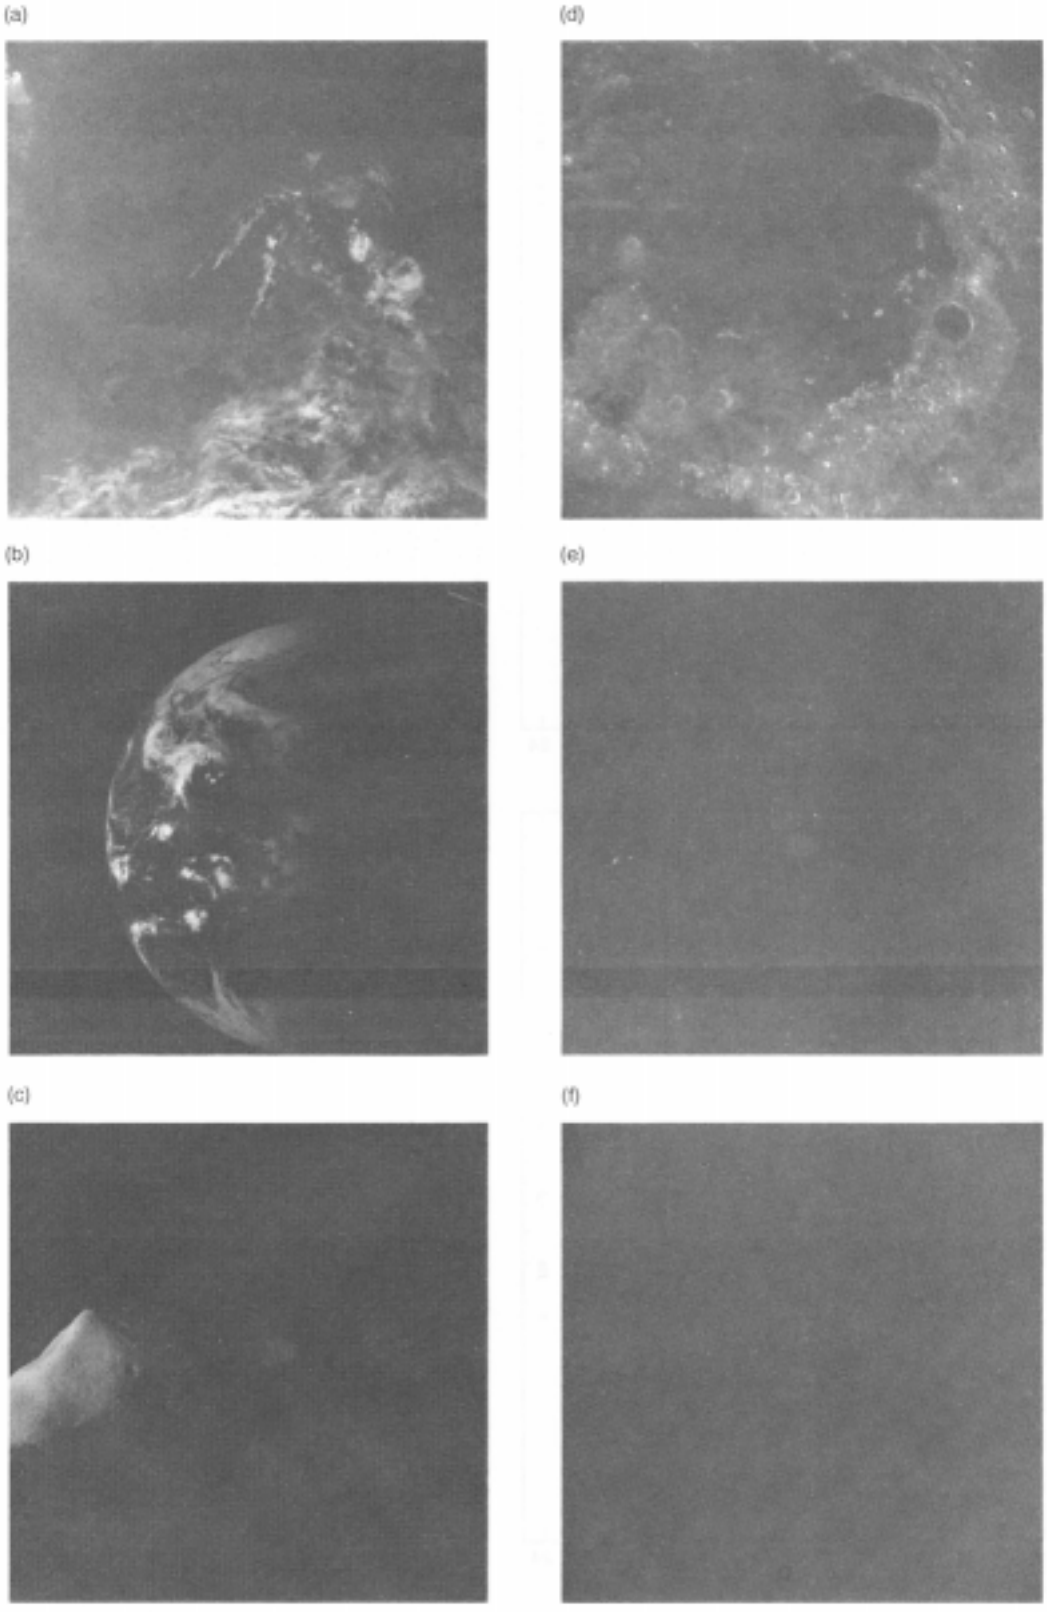

Fig. 1. Selected Galileo images: (a) rq534.ausvn, (b) rq534.ear9n, (c) rq534.gasn, (d) rq534.lunvn, (e) rq534.litn, and (f) rq534.pctn.

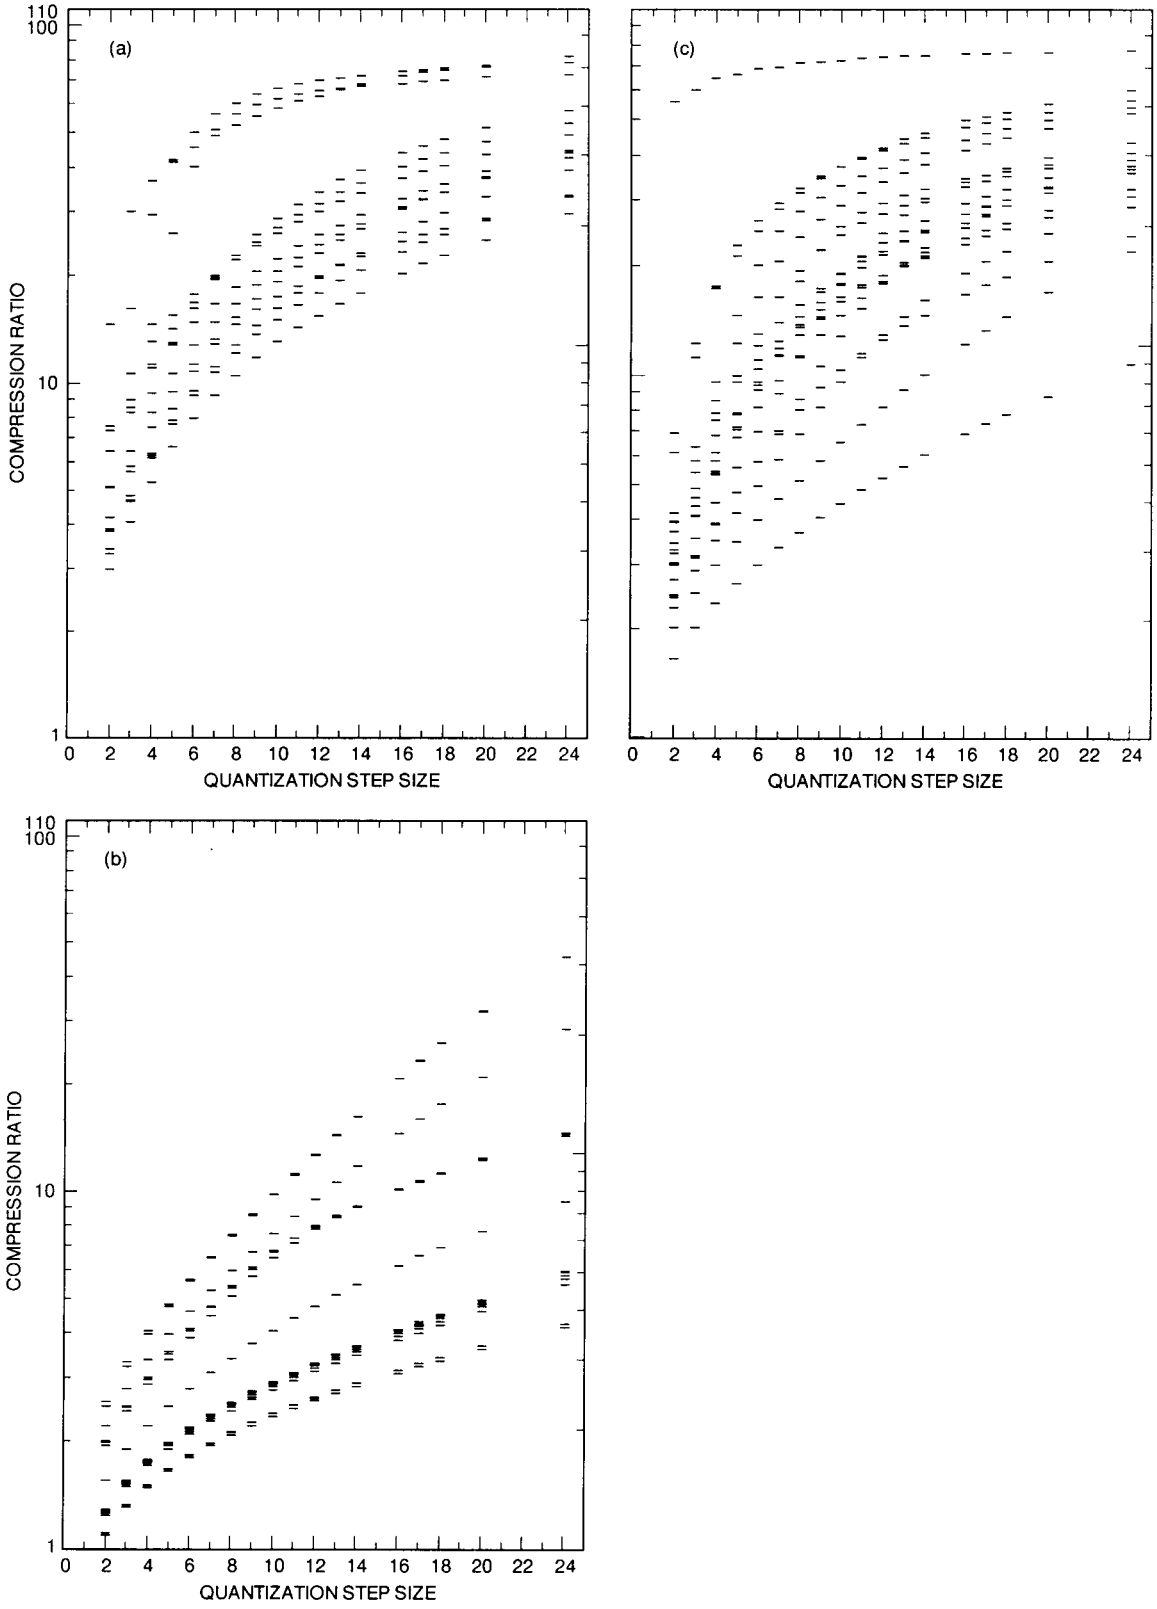

Fig. 2. Compression behavior for the small image sets: (a) compression ratio for selected Galileo SSI images, (b) compression ratio for images with radiation noise added, and (c) compression of simulated Galileo images from Voyager images.

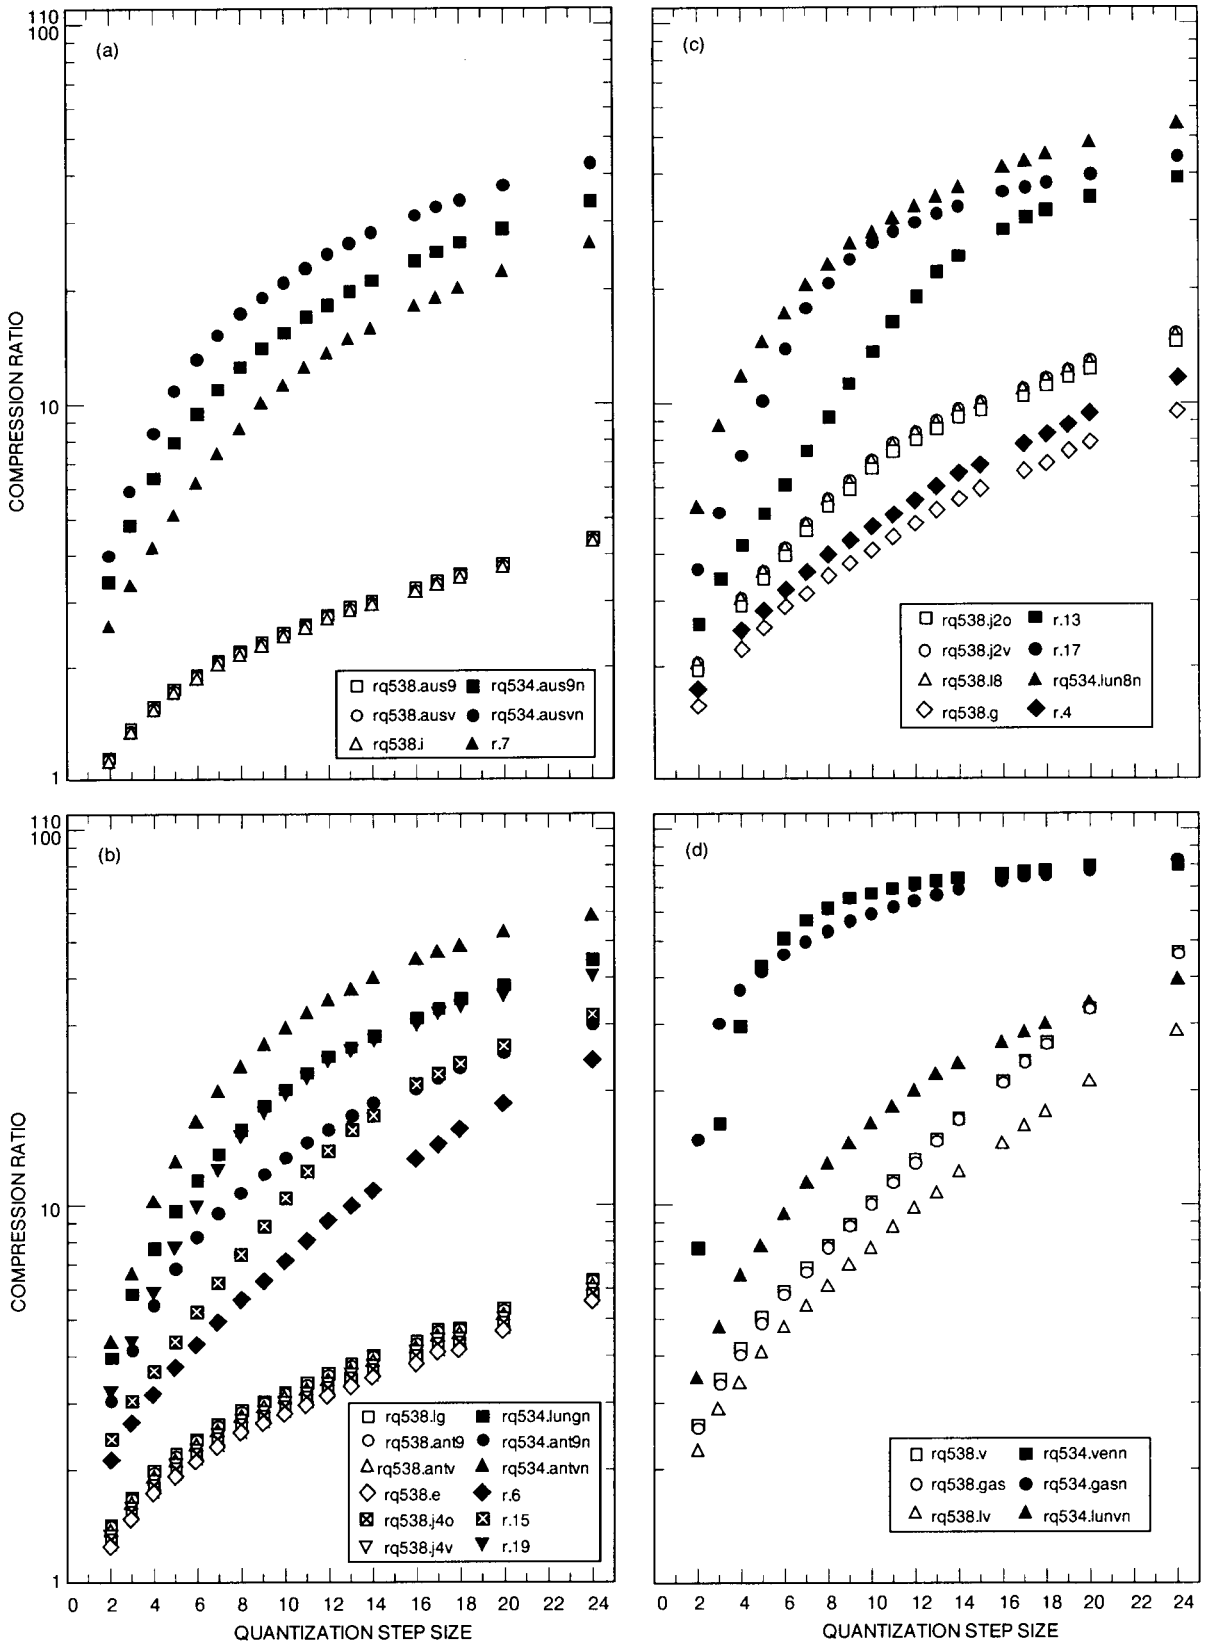

Fig. 3. Compression behavior for images with radiation noise added, showing the compression ratios for images with and without noise for (a) frame A, (b) frame B, (c) frame C, and (d) frame D.

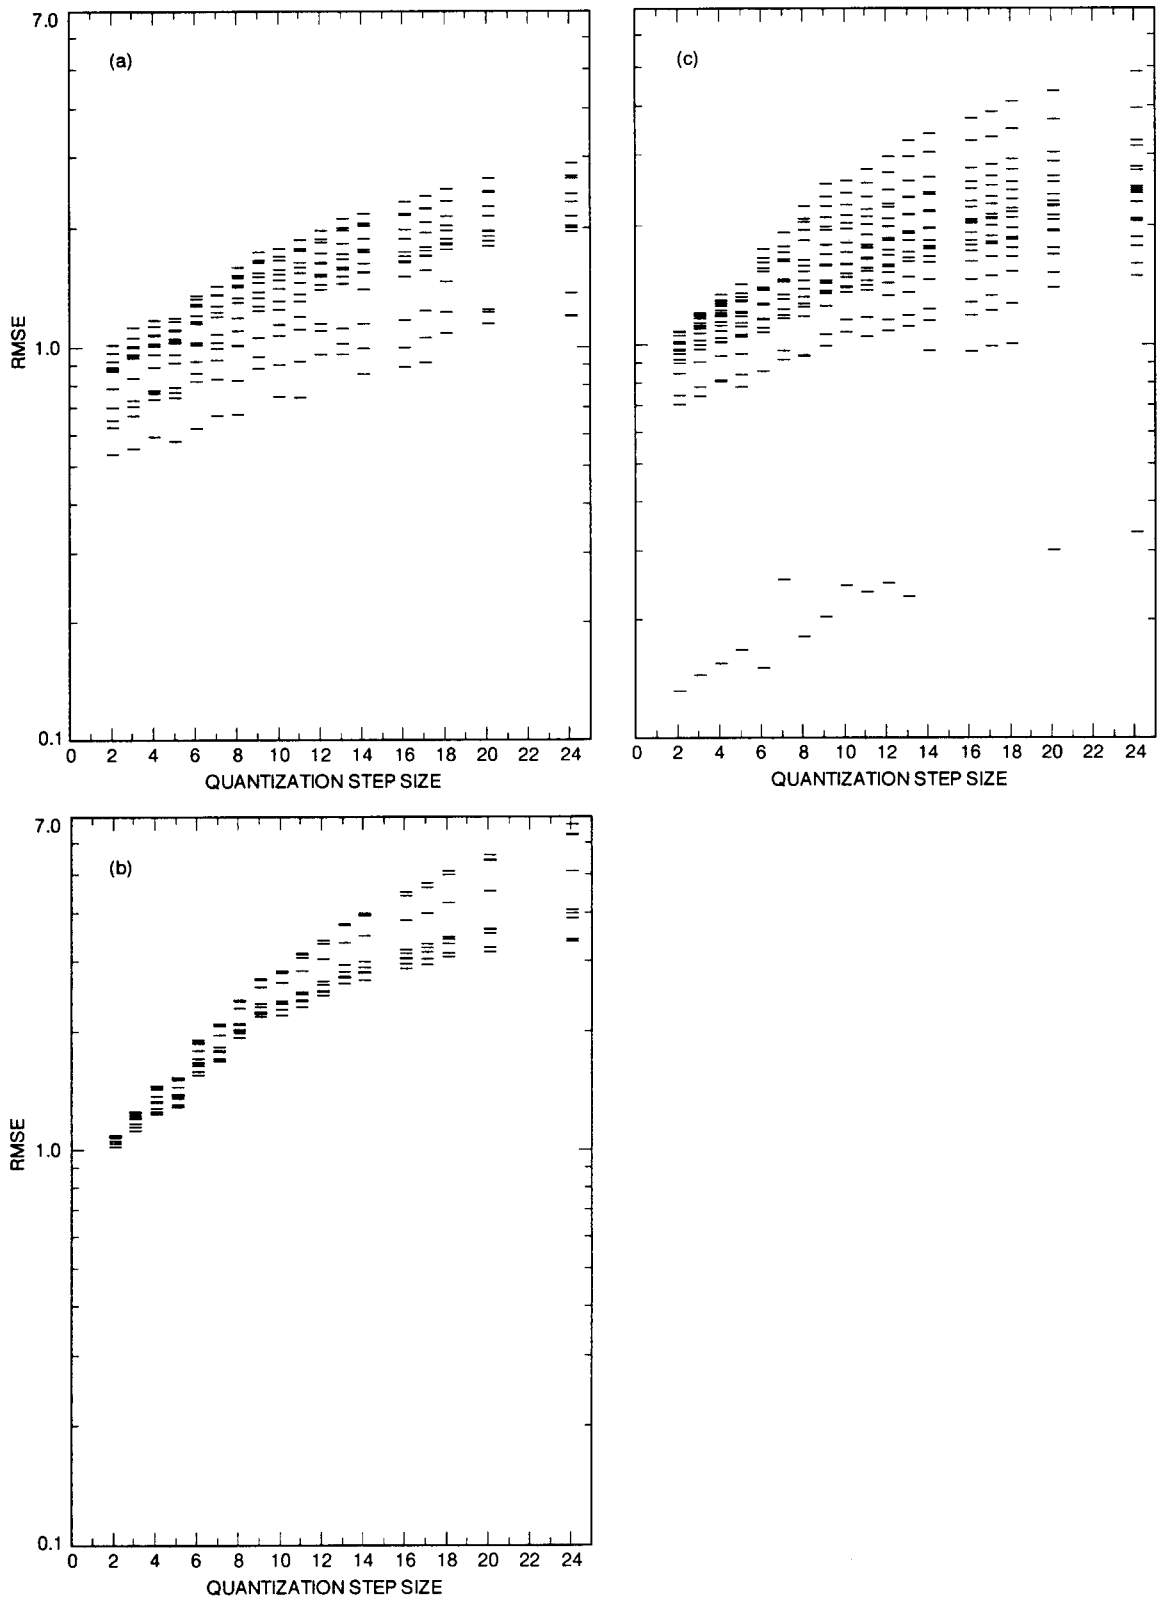

Fig. 4. RMSE distortion for the small image sets: (a) selected Galileo SSI images, (b) images with radiation noise added, and (c) simulated Galileo images from Voyager images.

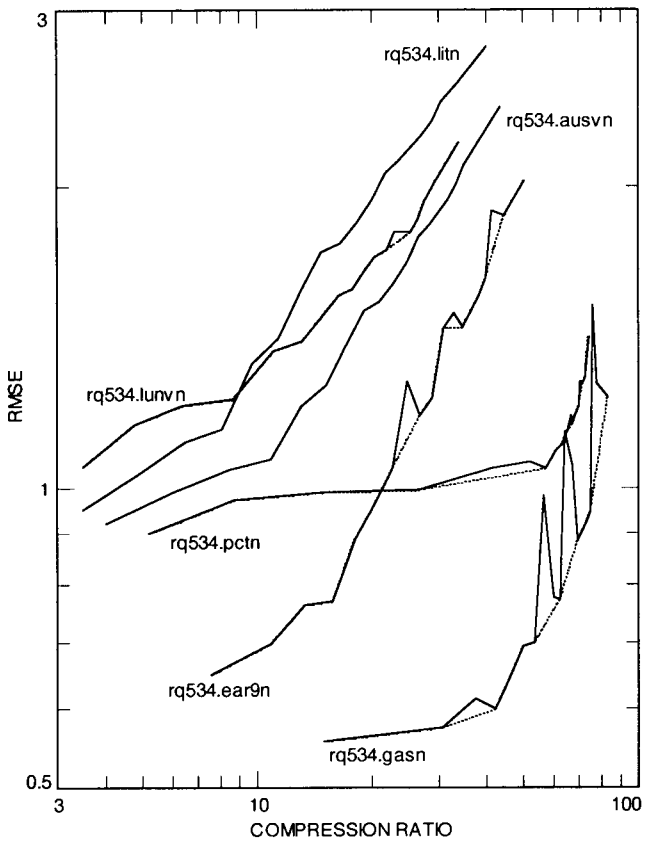

Fig. 5. Effects of inferior quantization step sizes on distortion versus compression curves.

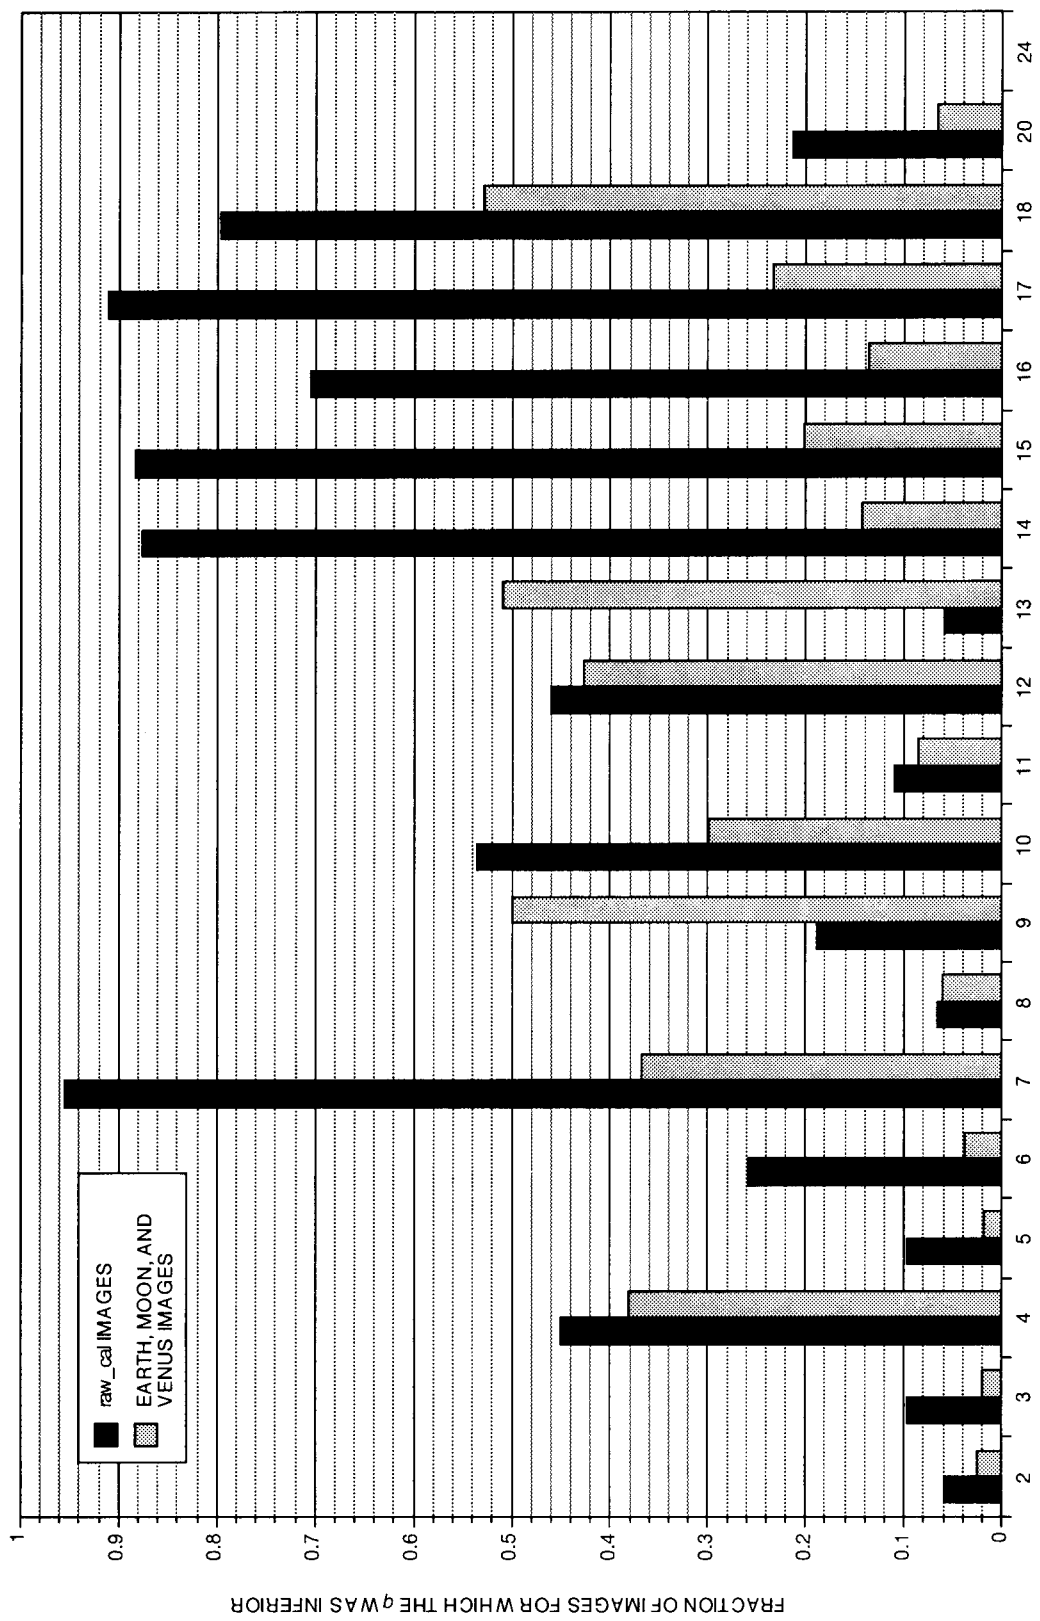

Fig. 6. Frequency of occurrence of inferior quantization step sizes over randomly selected images from the Earth 1 image set. 318 raw_cal images and 931 Earth, Moon, and Venus images were used to generate this histogram.

	filter mid-wave-length	*filter width*	*exposure time*	*gain*	*frame duration*
filter middle wavelength	1				
filter width	0.166	1			
exposure time	0.064	0.239	1		
gain	-0.32	-0.327	-0.1943	1	
frame duration	0.04	-0.228	-0.2214	0.025	1
mean	-0.04	0.196	0.62413	-0.08	-0.194
variance	0.02	0.032	-0.0763	0.1	-0.045
minimum	0.022	0.213	0.34753	-0.04	-0.595
maximum	0.071	0.214	0.09676	-0.22	0.0974
sample entropy	0	-0.266	-0.2223	0.023	0.2884
variance of differences	0.005	0.075	0.0184	-0.05	0.012
difference entropy	0.059	-0.045	-0.223	-0.23	0.2878
compression q=2	-0.04	0.151	0.25208	0.195	-0.359
compression q=3	-0.03	0.093	0.20209	0.145	-0.421
compression q=4	-0.03	0.089	0.18404	0.109	-0.433
compression q=5	-0.02	0.13	0.17664	0.062	-0.447
compression q=6	-0.01	0.201	0.17326	-0.01	-0.46
compression q=7	-0.07	0.223	0.14175	-0.06	-0.625
compression q=8	2E-04	0.282	0.16688	-0.1	-0.466
compression q=9	0.003	0.293	0.16243	-0.11	-0.466
compression q=10	0.003	0.301	0.15819	-0.13	-0.464
compressed filesize q=2	0.022	0.12	-0.1111	-0.24	0.1075
compressed filesize q=3	0.011	0.092	-0.0934	-0.17	0.1115
compressed filesize q=4	6E-04	0.016	-0.0938	-0.09	0.1278
compressed filesize q=5	-0.01	-0.058	-0.0951	-0.01	0.1396
compressed filesize q=6	-0.02	-0.097	-0.0956	0.035	0.1444
compressed filesize q=7	0.047	-0.026	-0.0809	0.025	0.2337
compressed filesize q=8	-0.02	-0.107	-0.0937	0.057	0.1425
compressed filesize q=9	-0.02	-0.101	-0.0914	0.057	0.1395
compressed filesize q=10	-0.01	-0.094	-0.0885	0.057	0.1345
rmse q=2	0.06	0.116	-0.2145	-0.36	0.1536
rmse q=3	0.063	0.125	-0.2009	-0.37	0.1604
rmse q=4	0.053	0.07	-0.2189	-0.33	0.1862
rmse q=5	0.047	0.073	-0.2123	-0.31	0.1888
rmse q=6	0.051	0.013	-0.2225	-0.27	0.2277
rmse q=7	0.085	0.06	-0.0063	-0.27	0.384
rmse q=8	0.014	-0.063	-0.2213	-0.17	0.2604
rmse q=9	0.024	-0.166	-0.2414	-0.13	0.3272
rmse q=10	0	-0.116	-0.2267	-0.11	0.2737

Fig. 7. A portion of the correlation matrix.

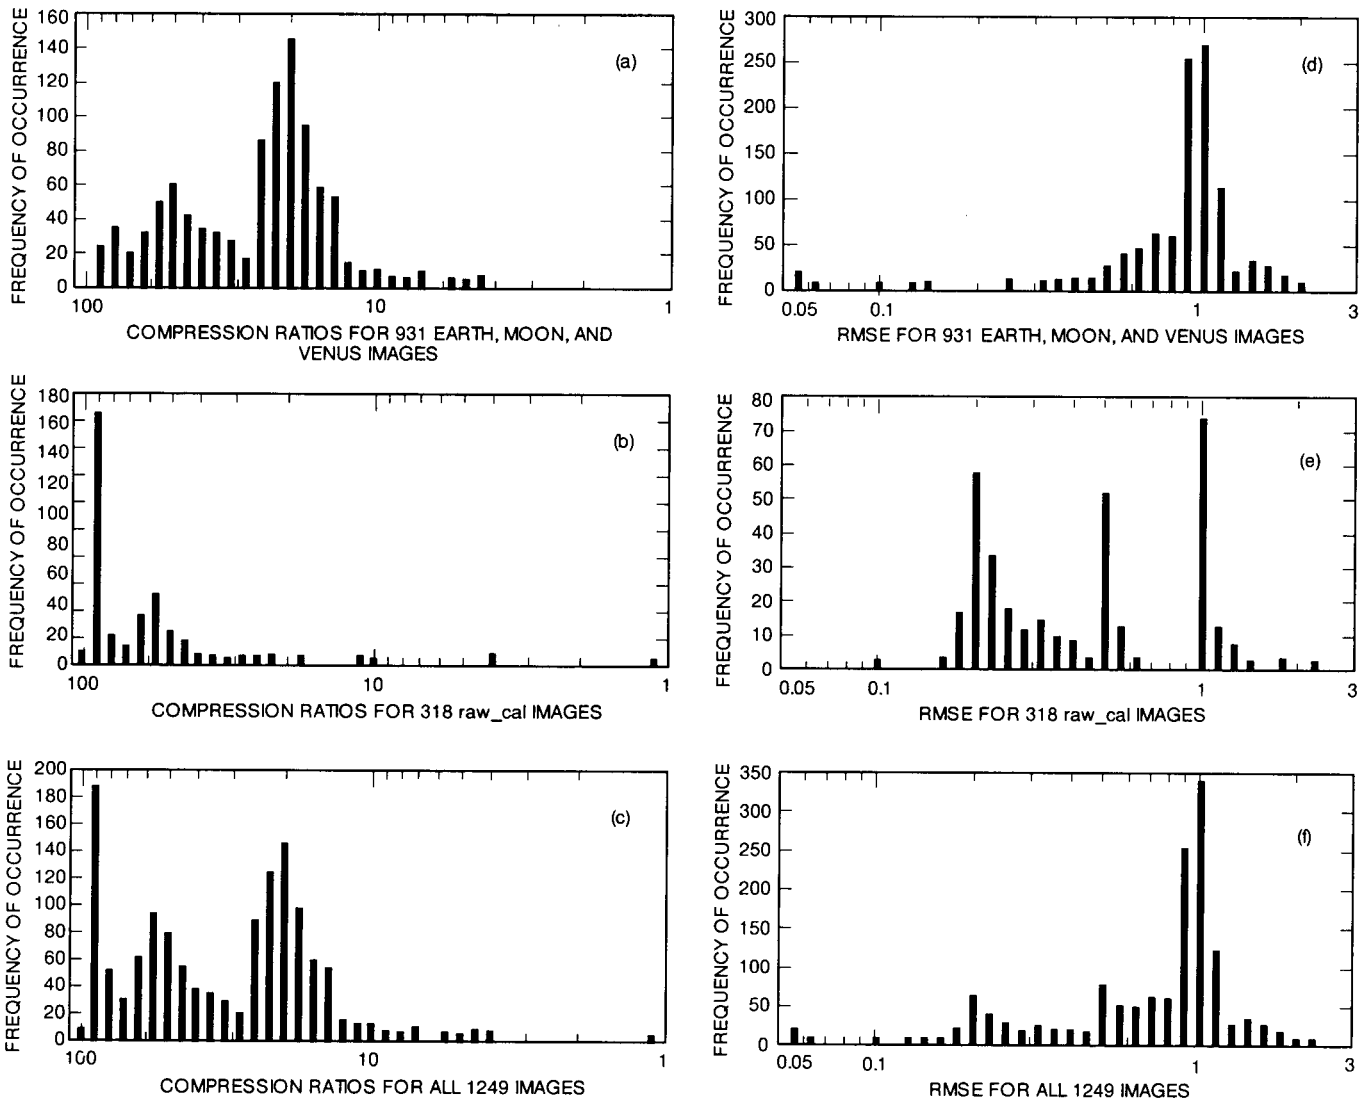

Fig. 8. Histograms of the compression ratios with quantization step size 8 for (a) 931 Earth, Moon, and Venus images, (b) 318 raw_cal images, and (c) all 1249 images, and of the rmse for (d) 931 Earth, Moon, and Venus images, (e) 318 raw_cal images, and (f) all 1249 images.

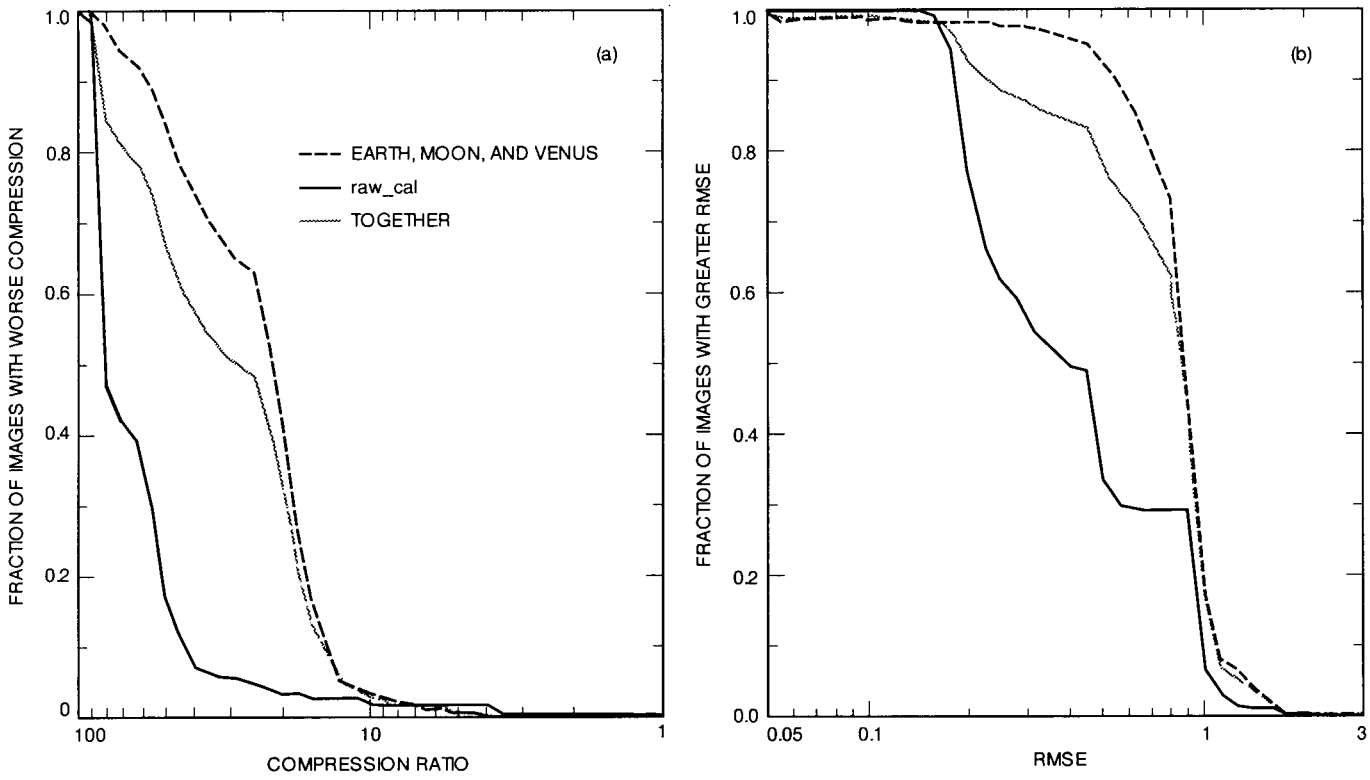

Fig. 9. Cumulative distributions for (a) compression ratios and (b) rmse for the randomly selected third of the Earth 1 images for quantization step size 8.