RIACS

Research Institute for Advanced Computer Science
NASA Ames Research Center

The Effect of Lossy Image Compression on Image Classification

Justin D. Paola

Robert A. Schowengerdt

RIACS Technical Report 95.18
August 16, 1995

Appeared in *Proceedings, 15th Annual IEEE International Geoscience and Remote Sensing Symposium*, Florence, Italy, July 20-14, 1995, pp. 118-120.

The Effect of Lossy Image Compression on Image Classification

Justin D. Paola

Robert A. Schowengerdt

The Research Institute for Advanced Computer Science
is operated by Universities Space Research Association,
The American City Building, Suite 212, Columbia, MD 21044, (410) 730-2656

Work reported herein was partially supported by the National Aeronautics and Space Administration under Contract NAS 2-13721 to the Universities Space Research Association (USRA) and under Grant NAG 5-2198 to the University of Arizona Department of Electrical and Computer Engineering. Work performed at the Research Institute for Advanced Computer Science (RIACS), NASA Ames Research Center, Moffet Field, CA 94035-1000

ABSTRACT

We have classified four different images, under various levels of JPEG compression, using the following classification algorithms: minimum-distance, maximum-likelihood, and neural network. The training site accuracy and percent difference from the original classification were tabulated for each image compression level, with maximum-likelihood showing the poorest results. In general, as compression ratio increased, the classification retained its overall appearance, but much of the pixel-to-pixel detail was eliminated. We also examined the effect of compression on spatial pattern detection using a neural network.

INTRODUCTION

With remote sensing studies becoming more global in nature, and computer processing power increasing, many scientists have been turning to larger and larger data sets. Unfortunately, storage of enormous data sets can be costly, thus making image compression an important consideration in the remote sensing field. For typical earth science imagery, lossless compression will result in about a 2:1 reduction. Lossy compression methods, however, commonly provide 10:1, 20:1, or even higher ratios, while maintaining the visual integrity of the image. The effect of these algorithms on supervised classification is important to consider before any data is archived with lossy compression.

JPEG IMAGE COMPRESSION

A common industry standard lossy compression method is JPEG (Joint Photographic Experts Group), which uses the discrete cosine transform. This algorithm is both fast and provides excellent energy compaction for highly correlated data [1].

JPEG makes use of the discrete cosine transform (DCT) for 8x8 contiguous sub-blocks of the image. The transform matrix $\mathbf{C} = \{c(k,n)\}$ is defined as:

$$c(k,n) = \begin{cases} \frac{1}{\sqrt{8}}, & k=0,\ 0 \le n \le 7 \\ \frac{1}{2}\cos\frac{\pi(2n+1)k}{16}, & 1 \le k \le 7,\ 0 \le n \le 7. \end{cases} \quad (1)$$

Most of the energy is packed into the first few transform coefficients. Varying levels of compression can be achieved by using variable quantization of these coefficients. Other compression algorithms, such as improved quantization of the DCT [2] and wavelet transform compression [3], are much superior both visually and in terms of mean square error, but are not yet image processing standards like JPEG.

EXPERIMENT

Recent studies have been reported on the effect of particular compression algorithms on subsequent multispectral analysis such as principal components and vegetation indexes [4] and on supervised and unsupervised classification [5].

In this experiment we have compressed four remotely-sensed multispectral images to varying degrees and have investigated the resulting supervised classifications obtained by the minimum-distance (MD), maximum-likelihood (ML), and three-layer

backpropagation neural network classifiers. We have also looked at the effect of compression on spatial pattern detection using a neural network.

The four classifications are:

- An urban land-use classification of Landsat Thematic Mapper (TM) satellite imagery of Tucson, Arizona, obtained April 1st, 1987.
- An urban land-use classification of TM imagery of Oakland, California, obtained August 1st, 1984.
- A geologic classification of Airborne Visible Near Infrared Imaging Spectrometer (AVIRIS) aircraft imagery of the Lunar Lake Volcanic Field in central Nevada, obtained September 29, 1989 [6].
- A combined temporal AVHRR NDVI (11 bi-weekly composites), spectral AVHRR, and DEM land cover classification of central California, using imagery from January to July, 1992.

For the first two images, the classifications were done two ways: 1) training on the original image with classification on the compressed imagery, and 2) both training and classification on the compressed imagery. For the second two images, all training was done on the original image. Both of the training methods are valid scenarios. In the first case, the user may have a few high quality (uncompressed) images to use for training, but desires to browse a compressed database. In the second case, the user is starting off with the compressed imagery.

The spatial data set is a series of synthetic aperture radar images from the Magellan spacecraft of the surface of Venus. A previous experiment on spatial pattern detection of impact craters [7] was reexamined after compression of the imagery.

RESULTS

Three different measures of classifier accuracy are presented in the tables. For each case, the accuracy of the training sites is given. If training was done on the original (uncompressed) image, this measure gives an indication of how much the compression has distorted the class exemplar regions. If training was done on the compressed image, this measure shows how well the classifier was able to describe the distorted training data.

The second measure is the accuracy of test sites that are independent of the training data. This is given for the Tucson image and helps show the generalization of the classification.

The third measure is the percentage of pixels in the classification of the compressed image (whether trained on the compressed image or not) that match the classification of the original image. It is safe to assume that the classification of the compressed data will be *no more* accurate than that of the original data. Thus, this measure gives a maximum bound on classification accuracy.

All of the classifications performed well for moderate compression ratios. In general, the maximum-likelihood and neural network classifications were more accurate on the original images than minimum-distance. The ML classifier, however, tended to deteriorate the most with increased compression. For the Tucson image, both the training and independent test sites degraded much more rapidly for ML than for the other two classifiers, as did the % match measure. Fig. 1 shows how the classifiers performed, after being trained on the original Oakland image, on a 28.5:1 compressed image.

It is intuitive that the MD classifier would not degrade as quickly as a parametric classifier. While individual pixel values can become quite distorted, and the class

distributions can change significantly with high JPEG compression (the classes tend to lose their spectral correlation, see [8]), the class means, on which the MD classifier depends, remain relatively constant.

Meanwhile, the assumption of a Gaussian class distribution, shaky to begin with in these types of classifications [9], causes the ML classifier significant problems as the pixel values change. The neural network, which derives a class distribution non-parametrically from the training data, suffers if the pixel values change significantly, but often starts off with a better description of these distributions and has more leeway for error.

As the compression ratio increases, the 8x8 image blocks become more homogeneous. The elimination of high frequency detail leads to a loss of detail in the resulting classification. Thus, while the overall classification remains fairly accurate, with large-scale spatial regions generally maintaining the correct classes, much of the finer detail is eliminated.

For the spatial pattern detection, the neural network windows were expanded to 25x25 and only one image band (Magellan SAR) was used. The compression in this case seemed to have little effect on the detection of impact craters.

CONCLUSION

Overall, it was found that high quality classifications could be obtained with any of the classifiers for JPEG compression ratios approaching 10:1 or even higher. Qualitatively, the classification retains its overall appearance, but the smoothing effect of high compression tends to eliminate much of the pixel-to-pixel detail. As expected, training on the compressed imagery could raise the training site accuracy, but did not raise the percentage of pixels matching the original classification. For the spatial pattern detection example presented, even severe image compression had little effect on detection ability.

REFERENCES

[1] A. K. Jain, *Fundamentals of Digital Image Processing*, Prentice-Hall, Englewood Cliffs, NJ., 1989.

[2] M. W. Marcellin, P. Sriram, and K.-L. Tong, "Transform coding of monochrome and color images using trellis coded quantization", *IEEE Transactions on Circuits and Systems for Video Technology*, pp. 270-276, Aug. 1993.

[3] J. M. Shapiro, "Embedded image coding using zerotrees of wavelet coefficients", *IEEE Transactions on Signal Processing*, pp. 3445-3462, Dec. 1993.

[4] S. S. Shen, J. E. Lindgren, and P. M. Payton, "Effects of Multispectral Compression on Machine Exploitation", *Proceedings, 27th Asilomar Conference on Signals, Systems and Computers*, IEEE 1058-6393/93, pp. 1352-1356, Pacific Grove, Ca., Nov. 1993.

[5] A. Habibi, B. Blyth, and C. Andrews, "Classification Consistency for Bandwidth Compressed Multispectral Imagery", *Proceedings, 27th Asilomar Conf. on Signals, Systems and Computers*, IEEE 1058-6393/93, pp. 1347-1351, Pacific Grove, Ca., Nov. 1993.

[6] E. Merényi, R. B. Singer, and W. H. Farrand, "Classification of the LCVF AVIRIS Test Site With a Kohonen Artificial Neural Network", *Contribution to the 4th Airborne Geoscience Workshop*, pp. 117-120, Washington, D.C., Oct. 25-29, 1993.

[7] J. D. Paola and R. A. Schowengerdt, "Comparisons of Neural Networks to Standard Techniques for Image Classification and Correlation", *Proceedings, 14th Annual IEEE International Geoscience and Remote Sensing Symposium (IGARSS '94)*, pp. 1404-1406, Pasadena, Ca., Aug. 1994.

[8] R. A. Schowengerdt and J. D. Paola, "Parallel Computing and Data Compression for Pattern Matching in Remote Sensing Image Databases", *Proceedings, Conference on Recent Advances in Remote Sensing, The European Symposium on Satellite Remote Sensing*, Rome, Italy, Sept. 26-30, 1994.

[9] J. D. Paola and R. A. Schowengerdt, "A Detailed Comparison of Backpropagation Neural Network and Maximum–Likelihood Classifiers for Urban Land Use Classification", to appear in *IEEE Transactions on Geoscience and Remote Sensing*, July 1995.

Figure 1: Classification of 28.5:1 JPEG compressed TM Oakland image.

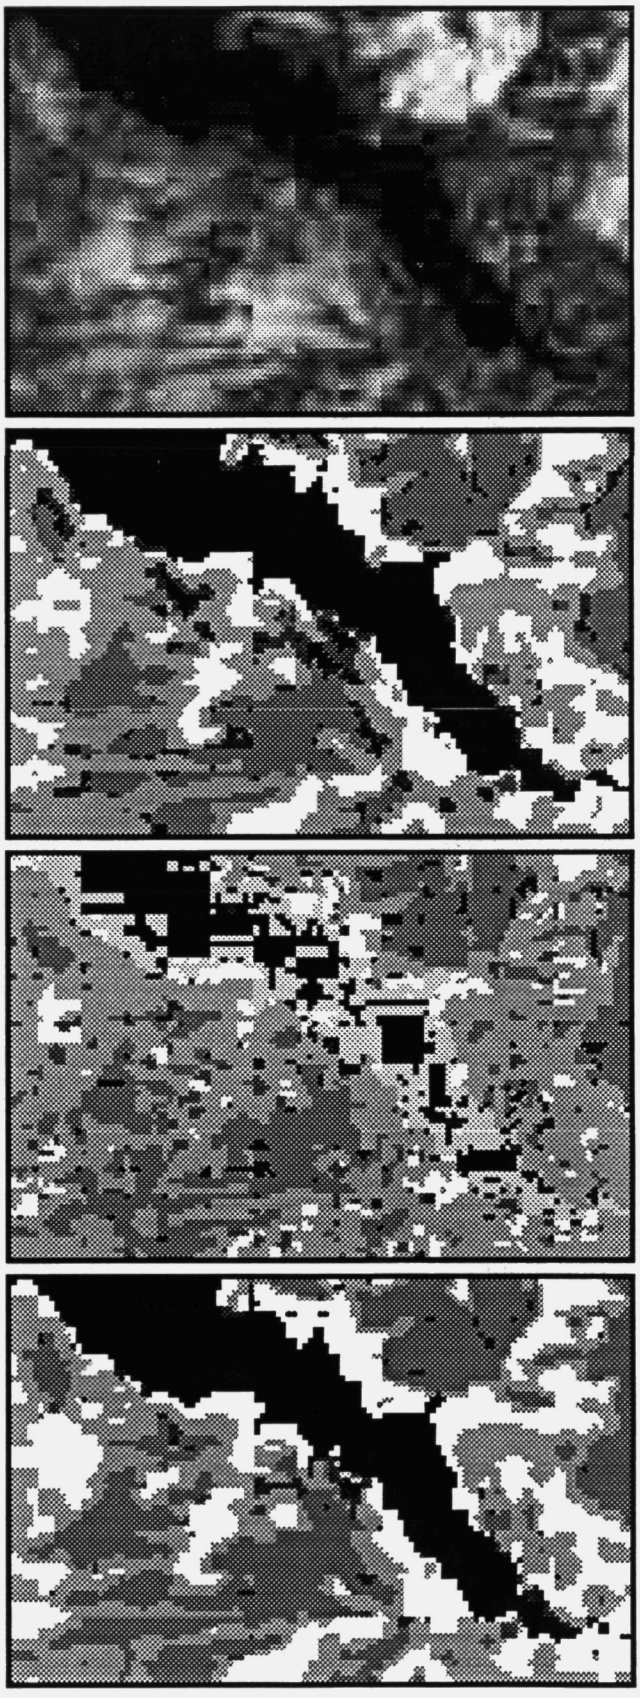

Band 5 of the compressed image.

Note: The 8x8 DCT blocking is apparent.

Note: The diagonal feature is a lake. White represents 'forest'. Other classes include 'grassland' and 'residential'. The class occupying much of the lake area in the ML map is 'urban'.

Minimum-distance classification

Maximum-likelihood classification

Neural network classification.

Table 1: Tucson TM image. Classifier performance on the compressed images when trained on the original data. Percent classification accuracies given for both the training sites and independent test sites. See Figure 2.

Image	Minimum-distance classifier accur. %	MD % match to orig.	Maximum-likelihood classifier accur. %	ML % match to orig.	Neural network classifier accur. %	Net % match to orig.
Orig.	Trn: 69.7 Test: 67.0	—	Trn: 96.9 Test: 89.5	—	Trn: 96.1 Test: 94.0	—
1.95:1	Trn: 69.5 Test: 67.2	97.1	Trn: 96.1 Test: 87.1	87.2	Trn: 95.5 Test: 93.6	95.3
3.8:1	Trn: 70.3 Test: 67.1	90.3	Trn: 86.3 Test: 77.9	71.9	Trn: 94.6 Test: 93.1	87.5
7.1:1	Trn: 69.2 Test: 67.4	83.9	Trn: 79.2 Test: 71.0	65.0	Trn: 93.9 Test: 92.3	82.3
16.9:1	Trn: 68.3 Test: 69.0	76.2	Trn: 65.4 Test: 61.3	54.6	Trn: 92.7 Test: 91.9	75.4
25.3:1	Trn: 69.2 Test: 69.4	72.4	Trn: 57.3 Test: 49.8	50.0	Trn: 92.4 Test: 91.0	71.8

Table 2: Tucson TM image. Classifier performance when both training and classification are carried out on the compressed images. See Figure 3.

Image	Minimum-distance classifier accur. %	MD % match to orig.	Maximum-likelihood classifier accur. %	ML % match to orig.	Neural network classifier accur. %	Net % match to orig.
1.95:1	Trn: 69.3 Test: 67.2	97.1	Trn: 96.4 Test: 91.6	86.3	Trn: 95.6 Test: 92.2	90.9
3.8:1	Trn: 70.4 Test: 67.2	90.2	Trn: 95.5 Test: 87.7	74.8	Trn: 95.1 Test: 91.3	84.9
7.1:1	Trn: 69.4 Test: 67.5	83.9	Trn: 96.2 Test: 85.8	65.1	Trn: 94.6 Test: 91.8	80.7
16.9:1	Trn: 68.0 Test: 68.1	76.1	Trn: 97.6 Test: 80.7	52.7	Trn: 95.7 Test: 91.6	73.3
25.3:1	Trn: 70.5 Test: 69.9	72.2	Trn: 96.4 Test: 72.6	46.7	Trn: 94.4 Test: 90.9	70.1

Table 3: Oakland TM image. Classifier performance on the compressed images when trained on the original data. Percent classification accuracies are given for the training sites. See Figure 4.

Image	Minimum-distance classifier accur. %	MD % match to orig.	Maximum-likelihood classifier accur. %	ML % match to orig.	Neural network classifier accur. %	Net % match to orig.
Orig.	89.1	—	97.1	—	95.3	—
1.65:1	89.1	99.4	97.1	98.9	95.2	99.4
3.1:1	88.9	94.7	95.1	91.3	94.5	95.1
5.3:1	88.4	90.5	93.6	86.4	93.9	91.2
13.45:1	91.2	82.9	90.9	76.4	94.6	85.1
28.5:1	93.9	76.0	83.9	66.8	93.8	78.8

Table 4: Oakland TM image. Classifier performance when both training and classification are carried out on the compressed images.

Image	Minimum-distance classifier accur. %	MD % match to orig.	Maximum-likelihood classifier accur. %	ML % match to orig.	Neural network classifier accur. %	Net % match to orig.
1.65:1	89.1	99.3	97.1	97.9	94.9	94.7
3.1:1	88.7	94.7	97.0	90.6	95.2	91.1
5.3:1	88.1	90.5	—	—	94.6	88.6
13.45:1	91.7	82.8	—	—	95.4	83.0
28.5:1	93.6	75.7	—	—	—	—

Table 5: Lunar Lake AVIRIS image. Training done on the original data only. Classification accuracies are given for the training sites. See Figure 5.

Image	Minimum-distance classifier accur. %	MD % match to orig.	Maximum-likelihood classifier accur. %	ML % match to orig.	Neural network classifier accur. %	Net % match to orig.
Orig.	88.6	—	100	—	98.1	—
1.45:1	83.7	91.2	88.3	84.9	89.6	88.3
2.45:1	80.7	81.7	72.8	66.0	79.3	75.8
3.5:1	74.8	72.4	66.1	51.8	73.4	63.8
7:1	28.3	35.4	11.7	15.6	13.1	5.3
12.5:1	27.2	25.1	13.5	13.9	10.1	10.1

Table 6: AVHRR NDVI time series/spectral/DEM image. Training done on the original data only. Classification accuracies are given for the training sites. See Figure 6.

Image	Minimum-distance classifier accur. %	MD % match to orig.	Maximum-likelihood classifier accur. %	ML % match to orig.	Neural network classifier accur. %	Net % match to orig.
Orig.	95.1	—	100	—	99.8	—
1.9:1	95.1	99.8	100	98.9	99.8	99.7
3.9:1	95.2	98.7	99.4	92.7	99.8	98.3
6.9:1	95.1	97.7	97.6	88.0	99.8	97.2
19.1:1	96.4	95.8	94.9	83.1	99.8	95.3
38.3:1	96.9	94.0	85.9	76.4	99.8	93.6

Table 7: Number of true (out of 11) and false crater detections in Magellan imagery of Venus using a neural network with 25x25 input nodes and 2 hidden layer nodes for various threshold levels (net output ranges from 0 to 1).

	Uncompressed image	5.9 : 1 compression	25.5 : 1 compression
Thr = 0.9	5/11, 0 false	5/11, 0 false	5/11, 0 false
Thr = 0.82	7/11, 2 false	8/11, 1 false	8/11, 0 false
Thr = 0.77	11/11, 4 false	11/11, 3 false	9/11, 4 false

Table 8: Same as above with training done on 14.5:1 compressed image.

	5.9 : 1 compression	25.5 : 1 compression
Thr = 0.9	6/11, 0 false	5/11, 0 false
Thr = 0.82	9/11, 2 false	9/11, 2 false
Thr = 0.77	11/11, 7 false	9/11, 8 false

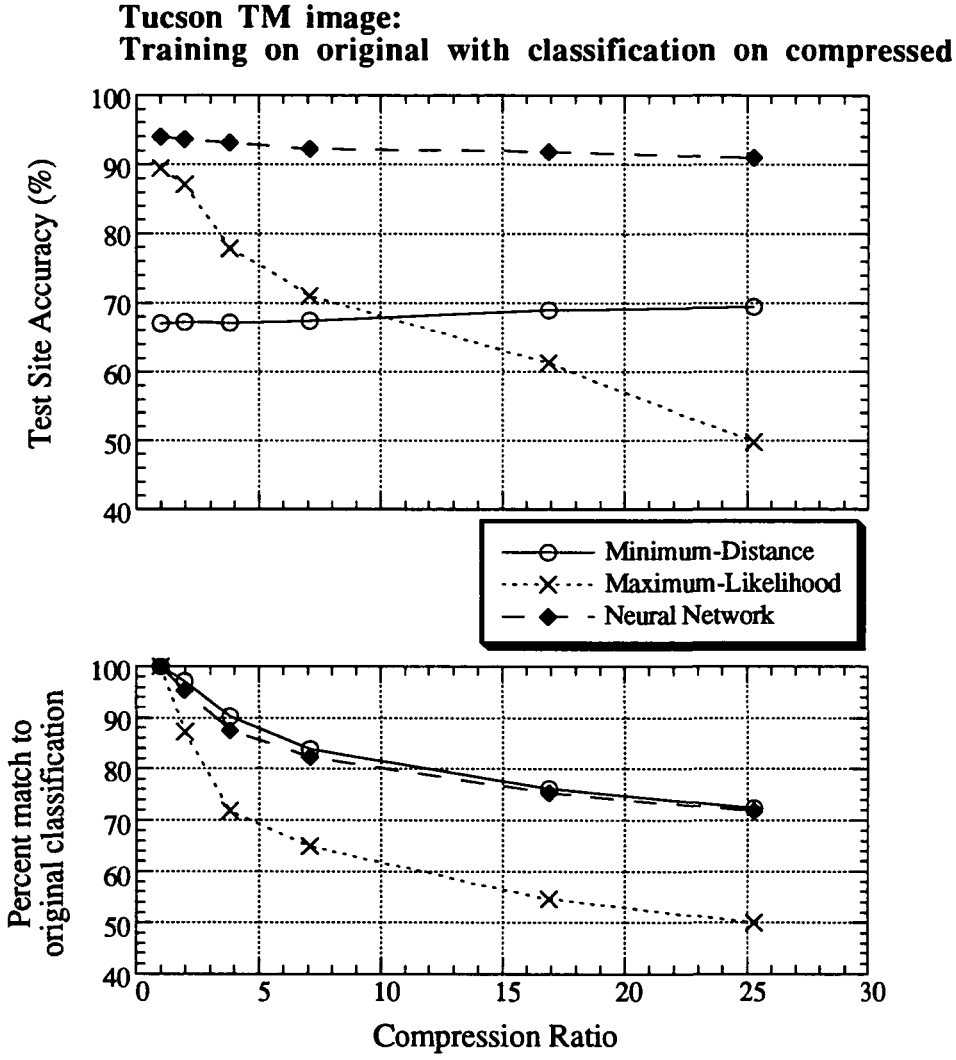

Figure 2: Tucson TM image. Classifier performance on the compressed images when trained on the original data. Classification accuracy is given for test sites independent of the training data. This data is from Table 1.

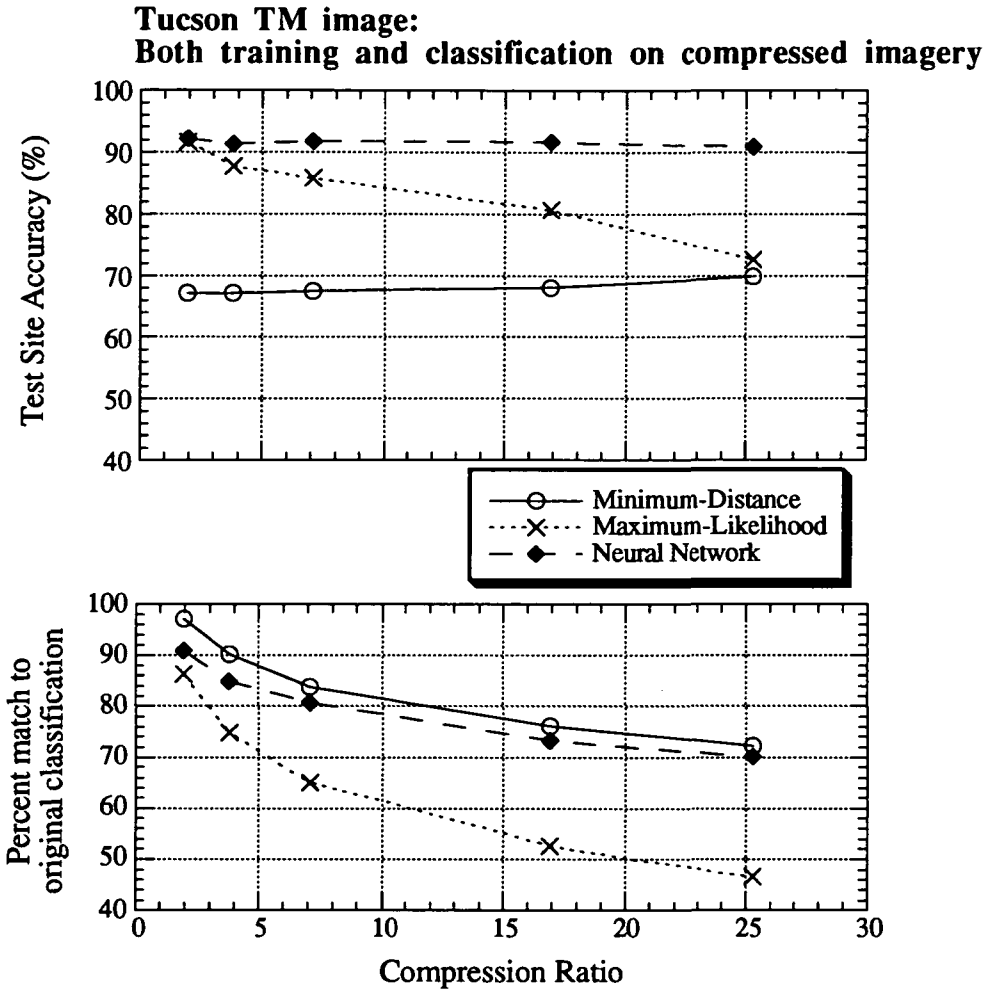

Figure 3: Tucson TM image. Classifier performance on the compressed images when trained on the compressed data. Classification accuracy is given for test sites independent of the training data. This data is from Table 2.

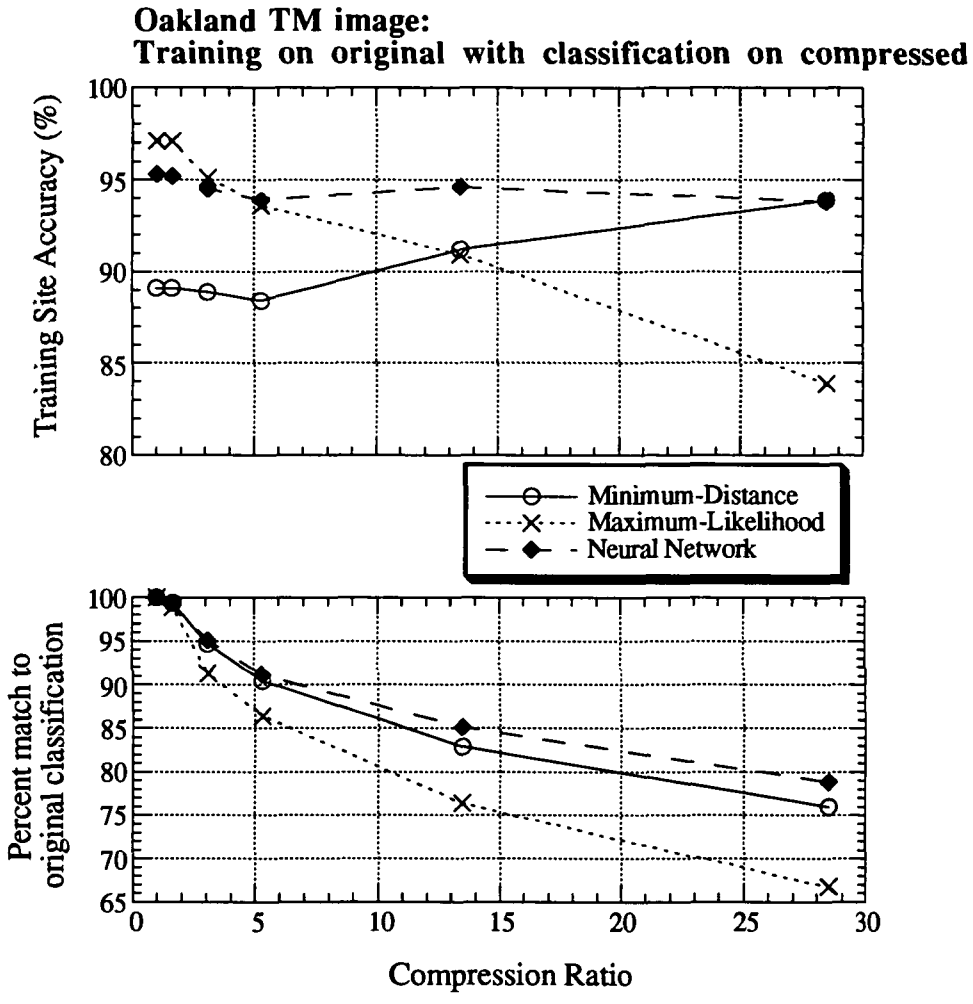

Figure 4: Oakland TM image. Classifier performance on the compressed images when trained on the original data. Classification accuracy is given for training sites. This data is from Table 3.

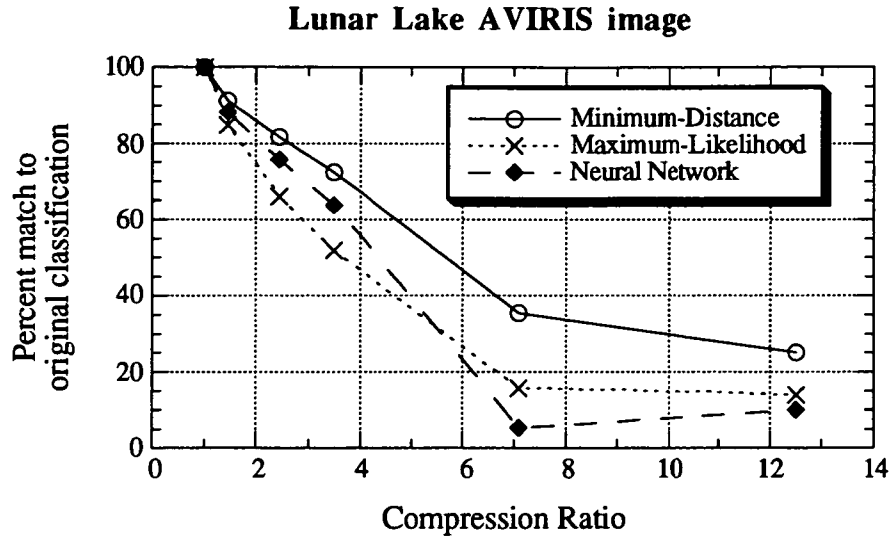

Figure 5: Lunar Lake AVIRIS image classifier performance. The data is from table 5.

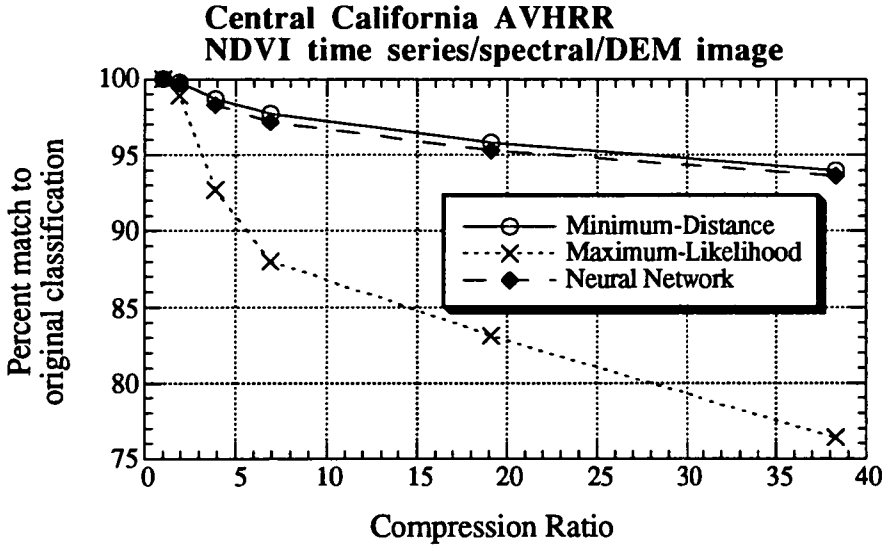

Figure 6: Northern California NDVI time series/spectral/DEM image classifier performance. The data is from table 6.

RIACS
Mail Stop T041-5
NASA Ames Research Center
Moffett Field, CA 94035