COMPRESSION OF DIGITAL COLOR IMAGES BY THE JPEG

Jussi Lammi and Tapani Sarjakoski

Finnish Geodetic Institute
Ilmalankatu 1A
SF-00240 Helsinki
Finland
Commission II

ABSTRACT

Image compression is not only desirable, it is also a necessity for the utilization of large digital images, e.g. digitized aerial color images. The JPEG standard proposal offers an alternative for carrying out the image compression task. The JPEG method itself and its suitability for photogrammetric work are studied, with special attention being paid to the geometric degradation of digital images in image compression. In our experience, JPEG image compression seems to be a good choice for color image compression. It gives a compression ratio of about 1:10 without considerable degradation in the visual or geometric quality of the image.

KEY WORDS: Image compression, standardization, geometric degradation.

1. INTRODUCTION

Compression of digital images is desirable because of the large volume of data in the images. It is especially useful in applications where many large images have to be archieved in a limited storage space or where digital images are transmitted over limited channels.

The basic idea in image compression is to remove redundancy from the image data. This is usually done by mapping the image to a set of coefficients. The resulting set is then quantized to a number of possible values which are coded by an appropriate encoding method. Nowadays, the most popular methods for removing redundancy are based on the discrete cosine transform (DCT), differential pulse code modulation (DPCM), vector quantization (VQ) and on the use of Laplacian pyramids.

The image compression standard proposed by the JPEG group offers a viable way of accomplishing the image compression task. This paper gives an overview of the JPEG compression method, which is still quite new to the photogrammetric community. Because the digital images used in the photogrammetric work tend to be very large, it is not reasonable to compress the whole image in a single compression step. A scheme for doing this is in smaller parts is proposed. The geometric effect of JPEG image compression on full-color images is empirically studied. We wanted to test our assumption that JPEG image compression does not affect the image geometry if small compression ratios are used. We also wanted to know what happens when the amount of compression becomes large.

2. IMAGE COMPRESSION BY THE JPEG

The following overview of the JPEG is based on an article by Wallace (1991), although other references also exist (C-Cube Microsystems, 1990; Storm Technology, 1990).

2.1 Background

The standard for image data compression discussed in this paper was proposed by the Joint Photographic Experts Group (JPEG), an ISO/CCITT working group, whose aim is to develop an international standard for continuous-tone still picture compression (ISO working group JTC1/SC2/WG10 in collaboration with CCITT SGVIII).

The JPEG traces its origins to videotext related project in the early 1980s. Its goal has, however, become more general as the number of applications needing a compression standard has grown. In the compression method selection process - conducted by the JPEG group - three compression algorithms out of 12 were chosen for a closer look in the 1987. Finally in early 1988 a discrete cosine transform (DCT)-based compression was chosen for standard development. In 1988 and 1990 the DCT-based method was further defined, tested and documented. After that the standard proposal started on its way through a normal ISO standardization process: from a Committee Draft (CD) to a Draft International Standard (DSI) and from there to an International Standard (IS).

The final ISO standard for image compression according to the JPEG will eventually be divided into two parts. Part 1 will specify the requirements and guidelines for the JPEG image compression, and Part 2 will contain the compliance tests taken. Although the JPEG standard proposal is already widely used in many applications, it will still be some time before the International Standard (IS) for the JPEG image compression is approved.

2.2 Baseline Sequential Encoding

The goal of the JPEG is to develop an image compression standard for compressing still-frame, continuous-tone images. The standard should cover as wide a range of applications as is feasible for a single standard. To satisfy the demand for versatile utilization the proposed JPEG standard is divided into four different modes: sequential encoding, progressive encoding, lossless encoding and hierarchical encoding. One or more distinct codecs (encoder/decoder -pairs) is specified for each of these modes.

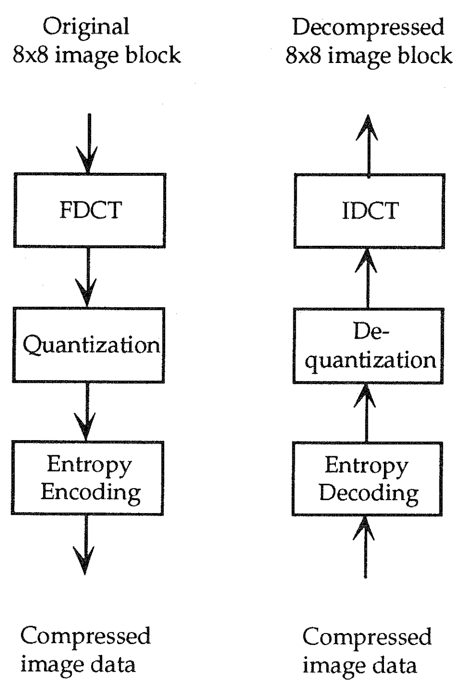

Figure 1. Baseline JPEG compression scheme.

Although the JPEG provides a variety of possibilities for encoding, it also gives a "basic" compression scheme - Baseline sequential coding - for straightforward use. In the following we shall concentrate on this Baseline method. This restriction is justified because Baseline encoding is sophisticated enough for many applications; it already explains the idea behind most of the JPEG modes and because totally lossless methods cannot provide sufficient data reduction when a huge amount of data is considered.

The Baseline coding method of the JPEG contains three sequential steps: Forward DCT (FDCT), quantization and entropy encoding (Fig. 1). The processing scheme is applied to a stream of 8x8 pixel blocks in grey scale images. The use of small image blocks takes into account the fact that the correlation between adjacent pixels is usually high in images of natural scenes. Decompression is achieved by following the processing steps in the opposite direction: entropy decoding, dequantization and Inverse DCT (IDCT).

2.2.1 FDCT For this step, the image is divided into 8x8 pixel blocks, each one of which is transformed by two-dimensional DCT, producing 64 output coefficients (set of basis-signal amplitudes). The transformed coefficient matrix is ordered so that the mean value of the coefficients (DC) is placed in the upper left corner of the grid; the remaining 63 coefficients (AC) are ordered so that low-frequency coefficients are closer to the DC value than the coefficients for high-frequencies.

2.2.2 Quantization After the FDCT, each cell in the 64 element coefficient matrix is quantized to a corresponding value in the predetermined quantization table. This is done to reduce the number of different coefficient values and to increase the number of zero value coefficients. Quantization is applied by dividing each DCT coefficient by the quantization step size and rounding the result of this division to the nearest integer. The quantized DC coefficient is encoded as the difference from the DC term of the previous block. All of the quantized coefficients (DC+ACs) are then ordered into a "zig-zag" sequence (Fig. 2). This makes entropy encoding more efficient because the nonzero low-frequency coefficients are placed before the high-frequency ones.

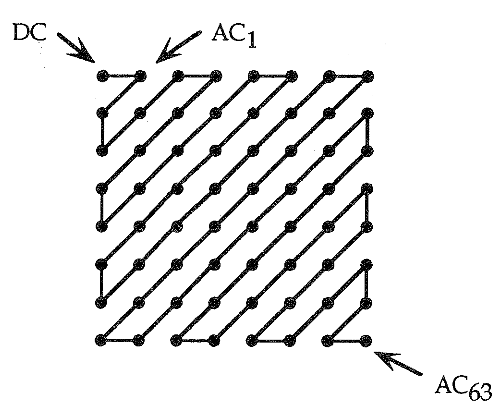

Figure 2. Zig-zag sequence.

2.2.3 Entropy Encoding The final step in the JPEG's DCT -based compression is entropy encoding, in which the quantized coefficients are losslessly encoded to a more compact form. Although the JPEG proposal specifies two entropy encoding methods (Huffman and arithmetic coding), the Baseline compression uses Huffman coding only.

2.2.4 Color Image Compression The Baseline sequential coding is for 8-bit images but it can be applied to color images as well. The color image compression can be done by compressing multiple channels one by one or by an approach in which the 8x8 blocks from each channel are compressed interleaved. Although the Baseline method compresses color images presented by any color model, it is best for images that are in color spaces such as YUV (Y for luminance, UV for chrominance) in which the color components are independent. Because the chrominance values need not to be considered as frequently as luminance values, the spatial resolution of the U and V components can be decreased. Subsampling of the chrominance channels reveals the advantage of using color spaces such as YUV in compression and explains why color images can be compressed with a better ratio than grey-scale images.

3. COMPRESSION OF LARGE DIGITAL IMAGES

In general, when large digital images are considered, it is not reasonable to compress the whole image in a single compression step. An approach in which the original image is divided into tiles "large enough" to be used as compression units is useful in applications in which only some of the image is required immediately. A catalogue indicating where each variable size block begins in the file is needed for quick access of image parts (Fig. 3).

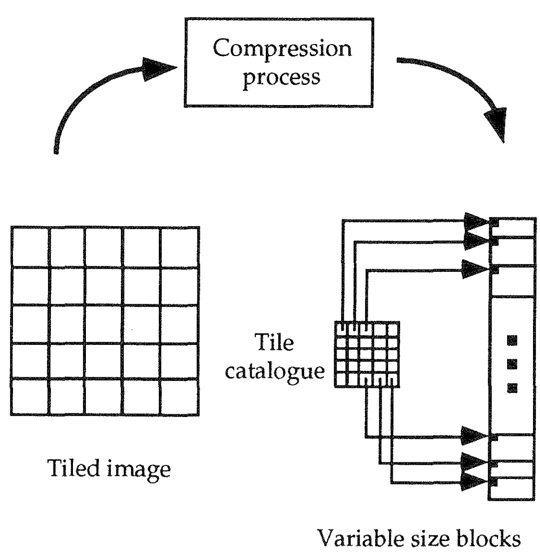

Figure 3. Piecewise compression of a large digital image.

In this piecewise compression scheme only those parts of the image that are really needed are decompressed. The time required for decompressing a few tiles is moderate compared with that for decompressing the whole image. Large images are usually also rather cumbersome to handle, and thus slow virtual memory swaps are necessary. Because the compression algorithms used for compression of large digital images are usually lossy, it is obvious that compression should not be done in intermediate steps.

4. GEOMETRIC DEGRADATION OF DIGITAL IMAGES IN THE COMPRESSION PROCESS

Geometric degradation of digital images in the compression process is especially interesting in photogrammetry because there the geometric quality of the images is important. The geometric quality of an image can be degraded in a compression process in two ways: 1) by degrading the radiometric quality of an image after compression such that measured targets are blurred or ambiguous and therefore pointing becomes inaccurate or impossible and 2) by shifting objects in line or column (or both) direction. Both of these degradations can be global or local by nature.

Geometric degradation depends to a considerable extent on the compression method used but it also depends on the contents of the compressed images and the amount of compression. It is clear that if the compression method is a lossless one then there is no danger of geometric degradation. On the other hand, all lossy methods are a possible source of geometric degradation of type one. Some of these methods might affect shifts as well - this applies particularly to methods where global knowledge is used for compression. Because compression methods usually have difficulties on spatial objects with large intensity variations in their neighborhood, it is reasonable to expect geometric degradation at these targets.

A research of (Akey & Mitchell, 1984) is an example of a study where geometric degradation of the compression process is examined. In their work, they studied the geometric effect of a discrete cosine transform-based image compression. It was found that the DCT-based compression moved the measured cross targets as much as 0.5 pixels when the compression ratio was about 1:16 on 8-bit images. As assumed, the error decreased when the compression decreased.

5. GEOMETRIC EFFECT OF THE JPEG

We studied empirically the geometric effect of the JPEG image compression on full-color images. We wanted to test our assumption that JPEG image compression does not affect to the image geometry if small compression ratios are used. We also wanted to know what happens when the amount of compression becomes large.

5.1 Test Arrangements

The test image was an aerial color photograph digitalized by the Sharp JX-600 desktop scanner in full-color mode and with a resolution of 600 dpi. This scanned image (Original) was compressed according to the Baseline JPEG image compression scheme, using the Storm Technology's PicturePress™ software with a Micron Xceed™ ICDP-II Picture Accelerator. The image was compressed into three different levels (Excellent, High and Fair), with compression ratios of about 1:7, 1:15, 1:66.

The visual quality of the image Excellent was very good. The image High was also quite good in the visual examination although some compression effects were seen. The visual quality of the image Fair was poor - the size of the compression blocks (8x8 pixels) was clearly visible, and all the edges were heavily smoothed. Despite this degradation in visual quality, we used the image Fair in the test as an example of a case in which the compression has become too large.

For the test 50 homogeneous linear features were chosen from the center area of the images. The features selected were larger than the compression blocks and were situated on areas where geometric degradation wass likely to occur (edges with large intensity variations). They were homogeneous in the sense that all were about 70 pixels long and represented the edges of roofs.

This test set was repeatedly measured on all test images using manual pointing - 20 times with the uncompressed image and 10 times with the compressed images. The linear features were pointed with visual interpolation. In this, sub-pixel pointing was achieved by zooming the viewed image area so that every pixel on the image was doubled before pointing.

The root-mean-square errors (rmse) for perpendicular differences between the endpoints of linear features were calculated. These root-mean-square errors were combined to represent the pointing precision of the linear features on each image.

To test the accuracy of the JPEG image compression, we kept a set of average features from the Original data as a reference when the rmse for perpendicular differences in the compressed test sets were calculated. The set of average features was formed by calculating averages from the endpoint coordinates of the linear features measured.

A set of average features was determined for each test set. Features from these sets were directly compared with the set of average features from the Original set. Again, root-mean-square errors for perpendicular differences between the endpoints were calculated.

5.2 Pointing Precision of Linear Features on Test Images

The pointing precision of the linear features on test images was calculated according to the test arrangements. Features with a perpendicular difference greater than 3 times the root-mean-square error of the corresponding data set were kept as gross errors and were rejected. The pointing precisions of the linear features on test images are presented in Table 1.

Compressed test sets were compared with the Original set one by one. The hypothesis H_0 in our case is written as

$$H_0:\ \delta^2_{\text{Original}} = \delta^2_{\text{Compressed}} \tag{1}$$

Table 1. Pointing precision of linear features on test images presented by maximum error (max v_y) and root-mean-square error (rmse). These values are expressed in pixels. Degrees of freedom for different test sets are also presented (υ).

Test sets (compression ratio)	max v_y	rmse	υ
Original	0.90	± 0.28	1842
Excellent (1:7)	0.84	± 0.27	886
High (1:15)	0.88	± 0.26	878
Fair (1:66)	1.15	± 0.39	864

and H_1 is

$$H_1:\ \delta^2_{\text{Original}} \neq \delta^2_{\text{Compressed}}\ . \tag{2}$$

The F-test was used to test the statistical significance of the hypothesis above. The ratio between the sample variances was used as a test parameter

$$z_F = \frac{s^2_{\text{Original}}}{s^2_{\text{Compressed}}}\ . \tag{3}$$

The hypothesis H_0 is rejected with the risk α if

$$z_F \geq F_{1-\alpha/2}(\upsilon_{\text{Original}}, \upsilon_{\text{Compressed}}) \tag{4}$$

or

$$z_F \leq \frac{1}{F_{1-\alpha/2}(\upsilon_{\text{Compressed}}\ , \upsilon_{\text{Original}})} \tag{5}$$

where $\upsilon_{\text{Original}}$, $\upsilon_{\text{Compressed}}$ are the degrees of freedom and F represents the F-distribution. The significance levels on which the appropriate null hypothesis H_0 can be rejected are presented in Table 2.

Table 2. The test set Original and other test sets are compared by using the F-distribution. The significance level on which the appropriate null hypothesis H_0 can be rejected is shown in percent. The test is based on the values presented in Table 1.

Excellent (1:7)	High (1:15)	Fair (1:66)
80.2 %	99.7 %	99.9 %

The pointing precision on the set Excellent does not differ from that on the set Original. These test sets are also considered equally distributed according to the F-test. But in the set High there is a difference between the compressed set and the set Original. However, the finding that the set High is even more precise than the set Original is somewhat surprising. One explanation is that the staircase effect on the test images is smoothed through image compression, and so the pointing is easier to repeat. The pointing precision in the set Fair is clearly poorer than the precision in the set Original. According to the F-test, the sets Fair and Original almost surely originated from a different data generation process. The compression used in this case is simply too large - the pointing of linear features is no longer unambiguous.

5.3 Pointing Accuracy of Linear Features on Test Images

The set of average features from the Original image were used as a reference when the rms errors for perpendicular differences were determined. The results from these calculations are presented in Table 3. The set of average features from the compressed images was also directly compared to the Original set of average features. Results from these comparisons are shown in the Table 4.

Table 3. The errors for perpendicular distances calculated by using the set of average features from the Original data. Maximum errors (max v_y) and root-mean-square errors (rmse) are presented. Values are expressed in pixels. Degrees of freedom for different test sets are also presented (υ).

Test sets (compression ratio)	max v_y	rmse	υ
Excellent (1:7)	1.00	± 0.30	886
High (1:15)	1.40	± 0.31	878
Fair (1:66)	3.32	± 0.65	864

Comparison of the results in the tables shows that the set Excellent is very close to the set Original. The set High is also quite close to Original. However, the values in Tables 3 and 4 indicate that there may be errors of a systematic nature in some features of High. In Fair the number of large errors has grown so much that the difference from Original is obvious.

After the test we checked the features that caused the largest residuals in the set Fair. The linear features having the four (4) largest misplacements in a perpendicular direction were caused by the fact that the image compression ratio was too large for those objects. The original edge had either disappeared totally or been smoothed so badly that the pointing was more like guesswork. In two of the cases the disappeared edge had a similar kind of linear feature in the close neighborhood. This coincidence caused a systematic error in the pointing.

Table 4. The difference between the set of average features from the test set Original and the set of average features from the compressed test sets. Maximum error (max v_y) and root-mean-square errors (rmse) are presented. Values are expressed in pixels. Degrees of freedom for different test sets are also presented (υ).

Test sets (compression ratio)	max v_y	rmse	υ
Excellent (1:7)	0.27	± 0.12	99
High (1:15)	0.73	± 0.18	99
Fair (1:66)	2.76	± 0.50	99

6. CONCLUSIONS

We have empirically studied the geometric effect of JPEG image compression on full-color images. The visual quality of the image Excellent (compression ratio 1:7) was very good, and no remarkable degradation in the geometric quality of this image was found in the test. The visual quality of the image High (1:15) had slightly deteriorated when compared with the Original one. A small geometric degradation effect in the case High was found in the test - some of the linear features were misplaced. In summary, Baseline JPEG image compression does not have a geometric effect on the image geometry when compression ratios of about 1:10 are used. Some geometric degradation effect may occur with higher compression rates.

Our examination was based on the use of visual pointing instead of numerical feature extraction methods. We believe that the results are applicable when digital images are used for visual, interactive mensuration in workstations. The use of JPEG image compression might have varying effects to different kinds of numerical feature extraction methods. This gives a topic for further research.

REFERENCES

Akey, M. L., Mitchell, O. R., 1984. Detection and Sub-Pixel Location of Objects in Digitized Aerial Imagery. A paper presented in the Seventh International Conference on Pattern Reciognition, Montreal, 1984.

C-Cube Microsystems, 1990. C-Cube Microsystems Workshop v 1.0.

Storm Technology, 1990. PicturePress User's Manual v 1.0.

Wallace, G., 1991. The JPEG Still Picture Compression Standard. Communications of the ACM. 34(4): 30-44.