Quality Assessment of the JPEG 2000 Compression Standard

Úlfar Steingrímsson and Klaus Simon
Swiss Federal Laboratories for Material Testing and Research (EMPA), St.Gallen

Abstract

Although the new JPEG 2000 data compression method has not yet truly reached the market, it has already achieved an unusually high level of attention in the industry. Since a new standard stands for major investments, JPEG 2000 has to offer convincing benefits for compensating the financial risks. In fact, compression rates of 1:200 or higher represent breakthroughs in image compression. In addition, the new data format supports a series of functions that are of major potential interest for software engineering. Ultimately, however, the visual attributes of JPEG 2000 are of utmost importance. While the visually perceived image quality is required to be still acceptable at highest compression rates for internet imaging purpose, it is at the same time expected to satisfy the topmost requirements at lower rates, particularly for the photography and the graphic arts industry. In an extensive empirical study, the question is investigated as to whether the high expectations of JPEG 2000 in respect to visual quality are justified. The paper gives an overview of the effected tests and presents the conclusions.

1. Introduction

In the end of 2001 the core coding system of the new JPEG 2000 image compression standard had been released and accepted as the ISO/IEC 15444-1 standard [1]. JPEG 2000 [2] includes a large list of ameliorations [3] with regard to the efficiency and the capabilities of the old JPEG format which is being widely used since the end of the 1980s [4–6]. JPEG 2000 offers new interesting features including substantially increased compression rates, progressive data transmission at increasing image quality, locally varied compression quality by using regions of interest, as well as the possibility of integrating security features (stamping, encryption, watermarking) [7]. Probably the most important innovation is the use of the Discrete Wavelet Transform (DWT) for lossless as well as lossy image compression in exchange for the Discrete Cosine Transform (DCT) as it is being used in JPEG.

Most important for JPEG 2000's acceptance will certainly be the visual quality of the compressed images. Since higher compression rates can only be achieved by accepting quality loss, only the direct comparison of JPEG 2000 and JPEG image qualities at different compression rates can give a reliable ranking of the new image format. A couple of studies actually indicate that JPEG 2000 seems to be favorable compared to JPEG [8–10]. The considered image collectives are however fairly small and the applied examinations only relying on numerical error measures such as the Peak Signal to Noise Ratio (PSNR) which do not necessarily correlate with perceptual image quality. Our study had been especially designed to balance this shortcoming and thus to provide visual quality examinations.

2. Methods

To achieve reliable quality analysis at varying compression rates, more than 14'000 interactive visual rating decisions had been taken by a total of 46 test persons. Special importance was either attached to highest statistical accuracy, or to a large image collection consisting of more than 140 test images (see fig. 1), alternatively. The image series had been evaluated by both, common computer users and skilled graphic arts specialists. Some fundamental questions have been investigated, namely:

1. Up to which compression rates can no visual deviations due to the compression process be discerned?
2. Which compression rates in JPEG or JPEG 2000 respectively lead to a comparable visual quality assessment?

The test series associated with these two questions are evaluated with the objective of ending up in a JPEG 2000 characteristic indicating at what compression rates and by what ratio the new image format can be expected to outperform the old JPEG.

The basis for the visual assessment was pursuant to print situation, i.e. the visual comparisons - between original and reproduction - were carried out in great detail and without time limits under standardized viewing conditions. The tests were conducted using the IBM T221 high-performance monitor and newspaper proof prints. For the

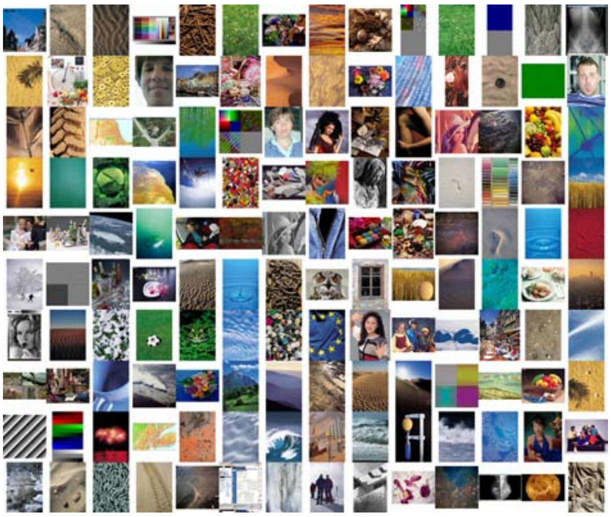

Figure 1: The high resolution image basis for the JPEG 2000 quality assessments.

JPEG compression, the corresponding modules from the JAI (Java Advanced Imaging) library were used. JPEG 2000 compression was done with the Java reference implementation of the standardization committee, obtainable from the EPFL [11]. Even though especially JPEG 2000 makes it possible to adjust lots of parameters for individual optimizations, it was largely the given standard settings that were used, since the average consumer is not expected to change them — except for the compression rates.

The JPEG 2000 study basically contains three types of evaluations designed to address the above questions. The first one, the discrimination test, is qualified for exploring up to what compression rates a compression technique is capable to produce images without visually perceivable deficiencies. The second one concerns about finding equal reproduction qualities for two different compression techniques at varying rates. The first two evaluations are optimized for high accuracy in the statistical sense, however being very time consuming and thus not suitable for large sets of images. The third one overcomes this drawback while basically addressing the same goals as the previous examinations.

2.1. Test 1 — Discernible compression rates

The test person was shown three different versions of a test image on the monitor, namely an original RGB-TIFF, a version compressed with JPEG or JPEG 2000 and subsequently decompressed, and a copy of one of the two. While for each test the positions of the different versions were randomly selected by the software, the task of the test person was to identify the image that was present only once. Correct and false identifications were documented by the software. The objective of this test was to establish the share of correct identifications per compression rate. Since the probability of random success is 33% after all, merely significantly higher hit rates imply correct identification. One single image comparison incident was granted about one minute of evaluation time. For statistical reasons, a total of more than 6000 tests had been performed by 46 different raters.

2.2. Test 2 — Comparable compression rates

In this experiment, the test viewer was shown a JPEG and a JPEG 2000 compressed image side by side on the monitor. He was then asked to indicate which image he considered to be the higher-quality version. During the test process the compression rates have been selectively varied by the test software to avoid arrangements of too dissimilar image qualities which do not yield meaningful ratings. An iterative self-learning algorithm was implemented to guarantee this request. For each image, the tests were carried on until the locations of equivalent JPEG / JPEG 2000 qualities had been reliably found within a broad compression range. In order to achieve a quality characteristic, a total of 200–300 comparisons were needed per image and test person requiring 3–5 hours, altogether.

2.3. Test 3 — Quality matching on large sets of images

To be capable to evaluate large sets of images the comparison process had to be dramatically accelerated. This intention was only feasible at the cost of statistical fidelity for each individual image review. This third type of experiment used a rather pragmatic mechanism for comparing the compression qualities of JPEG and JPEG 2000. Given a JPEG 2000 reference image at a certain compression rate, the quality of the opponent JPEG image was interactively varied by the examiner by adjusting a slider for manually increasing or reducing the JPEG compression rate, in so far as its quality begins to match the rank of the tag image. This process was repeated over the whole compression rates bandwidth.

3. Results

Fig. 2 highlights the overall distribution of the lowest compression rates being just discernible by eye from the original. The findings refer to the image collection outlined in fig. 1 which had been examined by two raters, each image. Surprisingly due to local structural loss, JPEG 2000 distortions were identifiable prior to JPEG artifacts, at increasing compression rates. Particularly regarding to the usage in graphic arts applications this finding implies that lossy JPEG 2000 is not necessarily the ultimate choice.

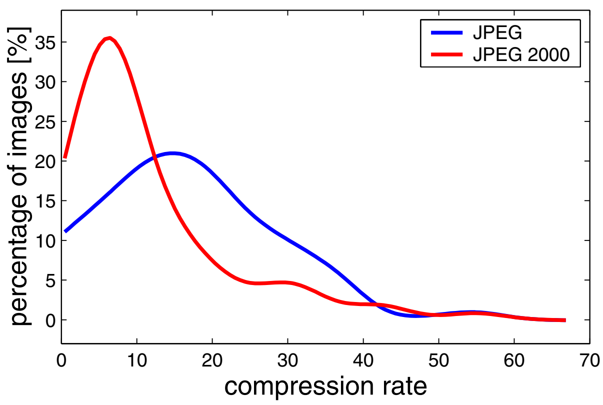

Figure 2: *Distribution of the maximal JPEG and JPEG 2000 compression rates being achievable without occurrence of perceivable image distortions (153 test images in total). Eye catching blurring of local structures cause early recognition of compression loss particularly in JPEG 2000 images.*

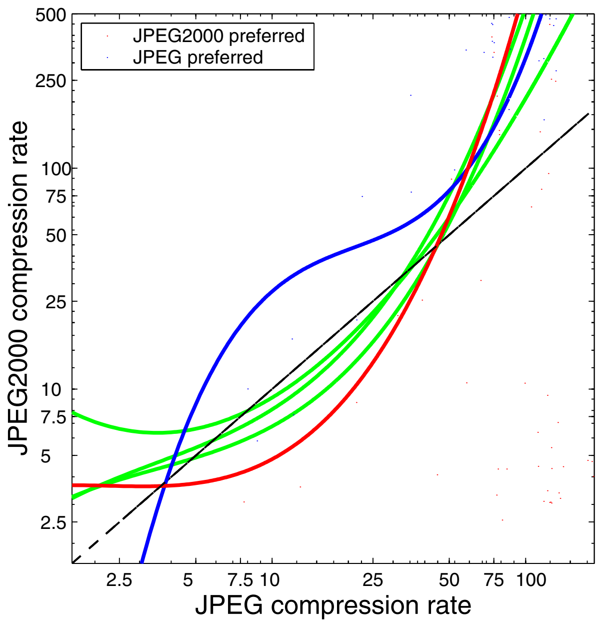

Figure 3: *JPEG and JPEG 2000 comparison tests carried out with different test persons and 5 different images. All lines show the average results for a single test image, each.*

The meticulous search for the separating line of equivalent compression qualities for JPEG and JPEG 2000 results in characteristics such as displayed in fig. 3. According to section 2.2, hundreds of JPEG / JPEG 2000 compression pairs had been visually compared by different examiners. Each single comparison was marked with colored points being either red or blue, meaning that either JPEG 2000 or JPEG compression had been favored. According to the total of all comparisons, the separating line between JPEG and JPEG 2000 preference was calculated. Fig. 3 includes the results of five different images. The *blue* line shows the evaluation for an image where JPEG 2000 is in most areas situated *above* the broken line of visual equivalent quality (*black* line), and hence produces clearly superior results compared to JPEG. In contrast, the *red* line shows the evaluation of an image, where JPEG produces considerably better results than JPEG 2000 at low and median compression rates. The characteristics of the remaining images (*green* lines) are located somewhere in between. As these examples clearly show, the motif-dependent variance is quite remarkable.

Special attention has to be payed to the compression rates lower than approximately 5. Since the quality loss is very low in this compression range and the deformations become almost invisible, the decision for either of the two compression techniques gets more and more randomized. Thus, a serious determination of the separating lines cannot be achieved at compression rates below 5.

The test outlined in section 2.3 was specially designed to find the line separating JPEG from JPEG 2000 quality for the large image collective in fig. 1. For this collection, the overall gain[1] of JPEG 2000 relative to JPEG is displayed in fig. 4. The graph denotes the median characteristics over the complete set of examined images and displays the gain of the JPEG 2000 compression rates within the full range of all rates that are feasible with JPEG. It is starling to notice that below compression rates of 1:50 the gain turns out to be negative. This means that for high quality compressions appearing at low compression rates JPEG often outperforms JPEG 2000 in terms of visually perceivable deficiencies. The variation between specific images is quite large, however. Moreover, the resulting image quality after a compression seems to be quite unpredictable. The results of the evaluations of large image sets (section 2.3) are well corresponding to those obtained from the tests designed for high accuracy (sections 2.1 and 2.2).

4. Argumentation and explanation

Visual investigations involve fairly intuitive global rating. In order to get a profound understanding on what is behind the above findings, a comparison of the particular background of the two compression techniques is required. Both, JPEG as well as JPEG 2000 lossy compression basically follow a workflow outlined below:

1. Partition of the original image into small blocks

[1]Gain in dB: $\varrho = 10 \cdot \log_{10}\left(\frac{\beta}{\alpha}\right)$

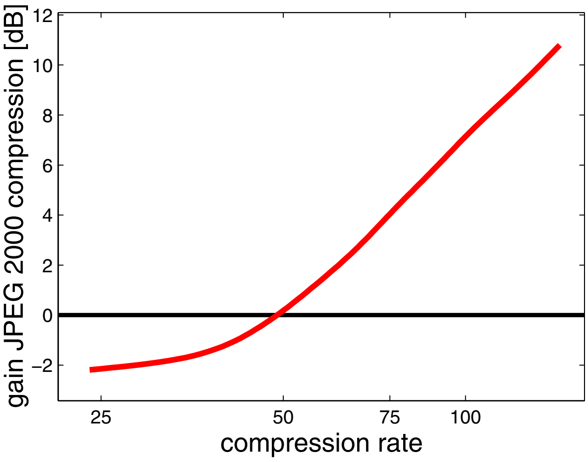

Figure 4: *Gain of JPEG 2000 at equivalent perceptible image qualities as JPEG. The graph shows the median characteristic over the grand total of 137 test images.*

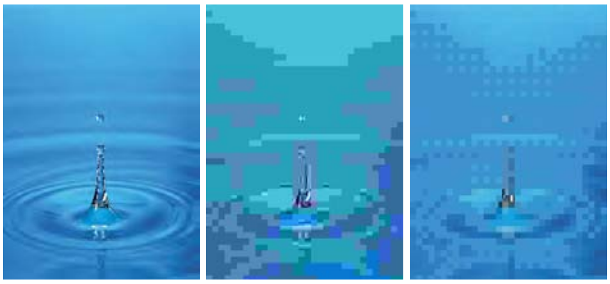

Figure 5: *Blocking artifacts are typical for JPEG (center image) and caused by the discontinuities due to the tiling of the original image into 8x8 pixel blocks. For JPEG 2000, tiling the image prior to compression causes heavy blocking artifacts, likewise (right image). If the image is however JPEG 2000 compressed without tiling as usual, the result (left image) is not distinguishable from the original. The compression rate is 1:120 for each of the three images.*

2. Orthogonal decomposition of each individual block
3. Removing non-essential information by quantization of the resulting coefficients
4. Lossless entropy coding of the remaining coefficients
5. Proper alignment of the coefficients for hierarchical image retrieval

While the actual image compression is effected in step 4, the compression related data loss is achieved with the quantization at step 3. Steps 1 and 2 are the preceding image transformations. They are substantially responsible for the distribution and the character of the distortions caused by the compression. The impact of these steps to image quality is investigated and presented separatedly.

4.1. Blocking artifacts

In JPEG, the original image is always tiled into small blocks of 8x8 pixels prior to compression (step 1). This tiling is performed for reasons of processing speed, since the calculation costs of discrete cosine transform (DCT) pertinent to JPEG increases with $O(n \cdot log(n))$. On the other hand, tiling results in discontinuities of the image function which is causing blocking artifacts. Blocking artifacts are typical for JPEG at heavy compression rates, especially for images with large-sized, soft color graduations. In JPEG 2000, the initial tiling process of the original image is feasible as well, however optional. Consequently without tiling, JPEG 2000 is capable to provide excellent results even at highest compression rates.

The left image in fig. 5 consists largely of a very low-frequency color background. JPEG produces heavy blocking artifacts already at moderate compression rates (fig. 5, center). In contrast with JPEG 2000, the untiled image is barely distorted even at high compression rates of more than 1:100. However, whenever initial tiling into 8x8 pixel blocks is achieved for JPEG 2000 compression all the same, the result shows distortions of at least the same impact as JPEG (fig. 5, right).

Consequently, the tiling is essentially responsible for the low image quality at high JPEG compression rates.

4.2. Orthogonal decomposition type

The most remarkable novelty of JPEG 2000 is clearly the use of the fancy discrete Wavelet transform (DWT) instead of DCT. The most important gain of DWT over DCT is in the first place the preservation of spatial information after the transform, and secondly, the possibility of individual tuning of the Wavelet shapes according to the specific requests. As a third advantage, the calculation cost increases only with $O(n)$. Except for the transformation speed, the benefit of DWT regarding to image compression is however debatable and had not been specially investigated in terms of perceivable image quality.

An estimate of this issue is achieved by a model experiment as it is exemplified in fig. 6. A simplified compression approach has been realized by transforming a picture into the frequency domain using either DCT or DWT, and by subsequently setting a certain percentage of the smallest coefficients to zero. The inverse transformation following reveals a picture with only a reduced amount of the original image information.

Fig. 6 illustrates the remainin reconstructing 100% down to only 0.025% of the original DCT / DWT coefficients. The DWT had been accomplished with the

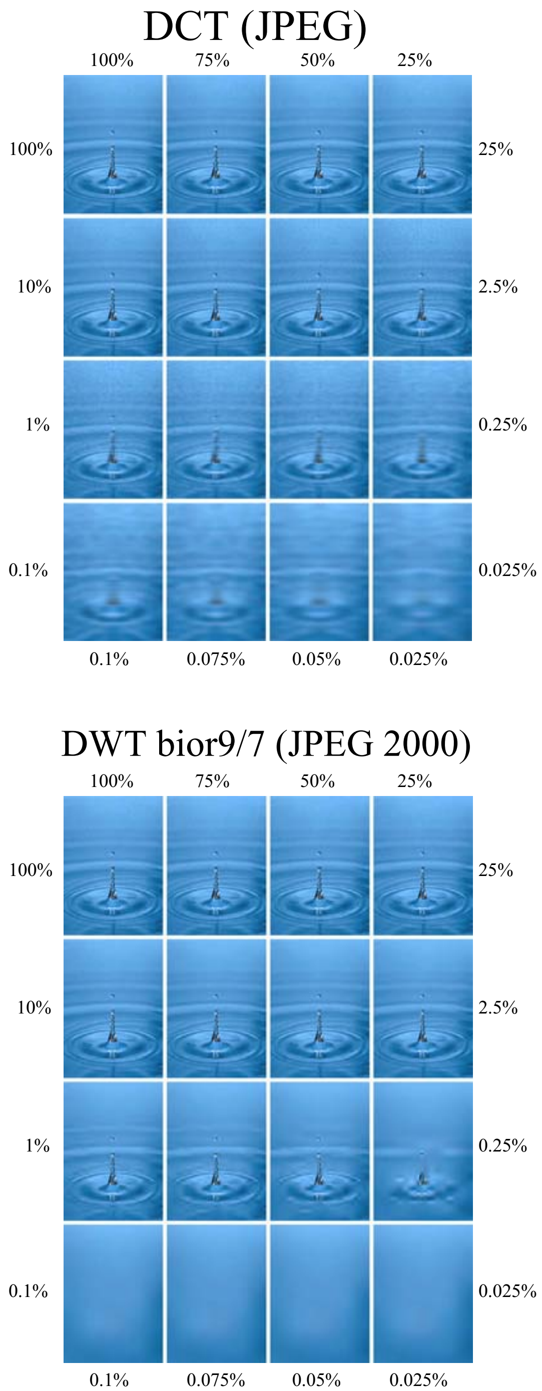

Figure 6: *Discrete cosine transform (top image) and Wavelet transform (bottom image), and following reconstruction with only a share of 100%, 75%, 50%, 25%, 10%, 7.5%, 5%, 2.5%, 1%, 0.75%, 0.5%, 0.25%, 0.1%, 0.075%, 0.05% and 0.025% of the largest DCT / DWT coefficients, while the remaining smaller coefficients were set at zero. Even with less than 1% of the original image information, the distortions in relation to the original image (top left) are acceptable for both transforms.*

biorthogonal 9/7 Wavelet employed in lossy JPEG 2000 compression. Both orthogonal transformations — DCT as well as DWT — reveal a rather similar perspective of being able to skip the required information to almost 1% until any distortions become perceivable. After more than 99% of the smallest coefficients are set to zero, the image quality decreases rapidly. The DWT shows no benefit compared to DCT at all. For the presented example actually the opposite holds true, albeit this statement is yet highly motif dependent.

From alike considerations can be concluded that DWT itself is generally not responsible for qualitative ameliorations in JPEG 2000. According to section 4.1, the real reason of the higher compressibility with JPEG 2000 is rather the omission of the initial image tiling.

4.3. Texture loss

The JPEG / JPEG 2000 comparison study presented in section 3 clearly averages a tendential benefit of JPEG at compression rates below 50. The essential reason for this finding was the loss of subtle structure in smooth image areas without outstanding edges (e.g. cloudy sky, sands).

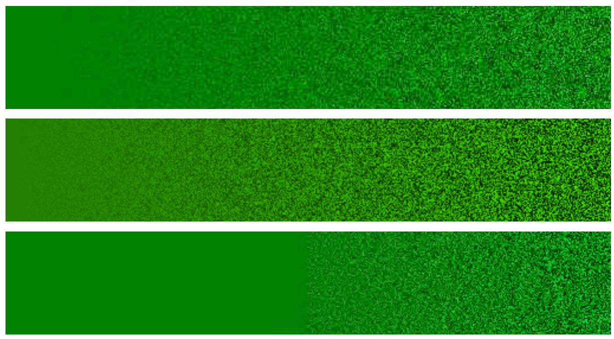

Figure 7: *White noise with an amplitude ranging from 0 (left) to 100% (only noise, right) is selected as the original image (center). The noise wedge is compressed at the compression rate of 1:24 JPEG (top image) and JPEG 2000 (bottom image). The JPEG image shows slight, uniformly distributed distortions, while in the case of the JPEG 2000 image the structural information disappears abruptly starting from the center of the image.*

Fig. 7 presents a step wedge with white noise, increasing from 0-100% from left to right. While the JPEG version (top image) shows a fairly well-balanced reproduction of the genuine picture (middle image) at a compression rate of 1:24 , the JPEG 2000 version at the same compression rate (bottom image) retains the structural information at high noise amplitudes only (right section), whereas beginning from the middle of the image a sudden loss of all structures occurs. The structural tear causes a considerably poorer overall image quality than the slightly distorted, uniformly distributed JPEG version. This example nicely illustrates a serious JPEG 2000 drawback.

In fig. 8 (left) we see a false color representation of a multiscale Wavelet decomposition of the original image from fig. 7 (center image) using the biorthogonal 9/7 Wavelet. The Wavelet coefficients, like the original image, consist of noise that increases uniformly from left to right. The right picture of fig. 8 shows the corresponding false color representation of the Wavelet decomposition of the JPEG 2000-compressed step wedge (fig. 7, bottom). A remarkable feature is that, after the decomposition stage, the left side of the high-frequency part of the image is empty. It becomes apparent that these coefficients are lost at the time of quantization.

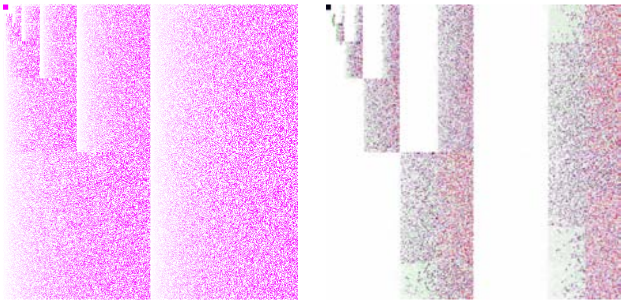

Figure 8: *False color representations of the Wavelet decompositions of the image from fig. 7 using the biorthogonal 9/7 Wavelet. The DWT of the original image is displayed on the left, the DWT of the JPEG 2000 version on the right. All small Wavelet coefficients were originally located in the left half of the image and totally eliminated by quantization.*

5. Conclusions

The total of all achieved evaluations yielded the following conclusions:

1. As expected, JPEG 2000 is capable to compress at highest rates far beyond the scope of JPEG. The perceptive quality comparisons moreover acknowledge JPEG 2000 being clearly superior to JPEG at compression rates above 50.
2. Most unexpectedly, the quality of JPEG 2000 underlies JPEG at compression rates below 50, on an average.
3. Those trends are highly depending on the specific image motif, however. At medium compression rates, the inter-image scattering turned out to be considerably higher than the differences between the two compression standards.

The findings from the evaluations providing high statistical accuracy basically agree with the results accomplished by the examinations at large image sets.

Further on, exclusive considerations on the image transformation aspects provided evidence that the skills of Wavelet transforms as used in JPEG 2000 is not superior to the DCT transform of JPEG in respect of to image compression. Better performance of JPEG 2000 at high compression rates rather turns out to occur because of omitting the initial step of tiling an image into 8x8 pixel blocks, as it is accomplished in JPEG.

Summing up, JPEG as well as lossy JPEG 2000 compression techniques are questionable for high quality imaging, however well suitable for medium and low quality applications where especially JPEG 2000 contains a big potential. Beyond doubt, the finding that JPEG 2000 might not necessarily yield better results than JPEG at lower up to medium compressions rates below 50 is fairly remarkable.

References

[1] Information technology - JPEG2000 image coding system - Part 1: Core coding system, December 2000.

[2] D.S. Taubman and M.W. Marcelling. *JPEG2000 Image Compression Fundamentals, Standards and Practice*. Kluwer Academic Publishers, Boston/Dordrecht/London, 2002.

[3] D. Santa-Cruz and T. Ebrahimi. An analytical study of JPEG 2000 functionalities. In *Proceedings of the IEEE International Conference on Image Processing (ICIP)*, volume 2, pages 49–52, Vancouver, Canada, Sep. 10-13 2000.

[4] Digital Compression and Coding of Continuous-tone Still Images, 1991. Draft International Standard / Commitee Draft 10918.

[5] W. Pennebaker and J. Mitchell. *JPEG Still Image Data Compression Standard*. Chapman & Hall, 1993.

[6] G.K. Wallace. The JPEG Still Picture Compression Standard. *IEEE Trans. Consumer Electronics*, 38(1):18–34, Feb. 1992.

[7] C. Christopoulos and T. Ebrahimi. The JPEG2000 Still Image Coding System: An Overview. *IEEE Transactions on Consumer Electronics*, 46(4):1103–1127, Nov. 2000.

[8] D. Santa-Cruz, T. Ebrahimi, J. Askelf, M. Larsson, and C. Cristopoulos. JPEG 2000 still image coding versus other standards. In *Applications of Digital Image Processing XXIII*, volume 4115, San Diego, California, Jul. 30 - Aug. 4 2000. Proceedings of the SPIE's 45th annual meeting.

[9] D. Santa-Cruz and T. Ebrahimi. A study of JPEG 2000 still image coding versus other standards. In *Proceedings of the X European Signal Processing Conference (EUSIPCO)*, volume 2, pages 673–676, Tampere, Finland, Sep. 5-8 2000.

[10] D. Santa Cruz and T. Ebrahimi. An Analytical Study of the JPEG2000 Functionalities. In *IEEE Int. Conf. Image Processing, Invited paper in Special Session on JPEG2000*, Vancouver, Canada, Sep. 2000. to be presented.

[11] R. Grosbois, D. Santa Cruz, J. Askelf, B. Berthelot, D. Bouchard, F. Henry, G. Mozelle, and P. Onno. JJ2000 - a JAVA implementation of JPEG 2000, version 4.1. internet site EPFL Lausanne, Mar. 2001. JJ2000 project by EPFL, ERICSSON and Canon Research Center France (CRF).