Perceptive Quality Estimations: JPEG 2000 versus JPEG

Úlfar Steingrímsson and Klaus Simon

Swiss Federal Laboratories for Material Testing and Research (EMPA), St. Gallen, Switzerland

The upcoming JPEG 2000 image compression standard is expected to replace the well-established JPEG format. It offers a wide range of data compressions, from lossless up to highest rates of lossy compression, in conjunction with a variety of useful new features which will presumably increase its popularity. Confident press reports suggest breakthroughs in image compression prospects. However, in the final analysis the most decisive factor in JPEG 2000's commercial acceptance will be its visual quality attributes. Accordingly, this article presents extensive estimations of the visually perceived image quality of JPEG 2000 compressions in comparison with JPEG. The study focuses exclusively on lossy data compression and its default settings. From the viewpoint of quality assessments commonly used in graphic arts industry, we present different types of interactive quality ratings applied to large data sets for characterizing image distortions at varying compression rates. In addition, the results are compared with PSNR considerations. The objective of this study is to arrive at an estimation of the potential benefits of JPEG 2000 relative to the JPEG standard with respect to compression quality, and in particular with regard to the requirements of the graphic arts industry.

Journal of Imaging Science and Technology 47: 572–585 (2003)

Introduction

In the late eighties, the ISO/IEC[1] committee introduced the JPEG[2] image compression standard,[1–5] allowing, in particular, heavy compression rates of up to 1:100, inducing a certain data loss, yet exhibiting acceptable quality degradations. Thus JPEG has currently been well established in areas featuring tight memory capabilities or low data transmission speeds, such as digital photography, internet or e-commerce applications.

The visual properties of JPEG are affected by the constraint of the basic image transformation to 8×8 pixel blocks.[3] Though not being of a fundamental type, this restriction was in fact induced by earlier hardware limitations and the increase in DCT processing time with $n \cdot log(n)$. Since the late eighties both technical and algorithmic knowledge — in particular with regard to the new concept of wavelet transforms — has augmented substantially. Owing to the boom in internet applications, requirements relating to image compression have increased dramatically in the meantime.

Work on JPEG 2000 started in 1997; the objective of incorporating modern image processing knowledge into a novel ISO/IEC 15444 image compression standard was explained in a couple of recapitulatory articles.[6–9] Detailed information is given in Ref. 10. We refer specifically to the core coding system part 1 (ISO/IEC 15444-1) which was released in 2001.[11] The standard has been and continues to be elaborated in cooperation with scientific with industrial parties. JPEG 2000 provides a long list of improvements and interesting new features.[12] Outstanding enhancements are[7] the possibility of applying a substantially broader bandwidth of compression rates, data transmission which is progressive by pixel accuracy or by pixel resolution, alternatively,

- a large color bandwidth enabling bi-level up to continuous tone compression,
- high robustness in the presence of bit errors,[13]
- the facility to provide variable data fidelity at specific image sectors using regions of interest (ROI),[14]
- support of image-related meta data,
- an open data architecture allowing easy adaptation of new image classes,
- facilities to provide protective image security features (watermarking, labeling, stamping, encryption),
- and, finally, backward compatibility of the marker and marker segment syntax with JPEG.[11]

Of the various new facilities, however, JPEG 2000's most celebrated innovation is undoubtedly the use of a wavelet transform instead of a Fourier-related transform for lossless and lossy image compression. This novelty is commonly supposed to be the major key to a significant improvement. There are indeed some studies indicating that JPEG 2000 provides a better image quality than JPEG,[15–17] even though none of these explorations explicitly identifies the wavelet transform as the basis of the presumed quality enhancement. Furthermore, most of the earlier investigations rely on PSNR (Peak Signal to Noise Ratio) as a quality measure. Whether or not the PSNR correlates with the perceived image quality is a controversial issue.[15]

Original manuscript received November 22, 2002

Color Plates 18–30 are printed in the color plate section of this issue, pp. 586–603.

The key question, however, is that of compression quality. In contrast to internet and TV applications, where the demand for high compression rates is of utmost importance, other business sectors such as photography and the graphic arts industry accept only high quality image compression causing invisible distortions, at most. For that reason, JPEG did not really become popular in graphic workflow:

- Even though the quality loss of JPEG images is supposedly imperceptible at low compression rates, slight shortcomings can be observed. High compression rates even lead to serious deficiencies typical of JPEG, such as tiling (**Color Plate 18, p. 594**) and color shifts (**Color Plate 19, p. 594**).
- Industrial processes involve repeated image compression and reconstruction cycles. Using lossy JPEG compression runs the risk of progressively losing image information and thus successively reducing image quality.

Hence, the acceptance of lossy image compression might dramatically increase in connection with the new format, provided that a substantially improved compression quality of JPEG 2000 can be verified at all compression levels. This issue is to be further investigated.

Transform-Based Lossy Image Compression

The basic idea of lossy image compression is the removal of redundant data due to pixel coding, interpixel, or psycho-visual redundancy.[18] It is achieved by first transforming the image values from the spatial domain into a statistically uncorrelated data set followed by consecutively removing the least relevant layers in the data coefficients. The image transformation is usually an orthogonal decomposition (Fourier, Chebyshev, Hermit, etc.) yielding a sparse representation of the original data, causing most of the transform coefficients to become zero or very small. Removing non-essential information is carried out by their subsequent quantization (or even elimination), this being the cause of data loss and thus of unrecoverable distortions.[10,19] The following entropy coding of the remaining coefficients is lossless and therefore affects compression size, though not compression quality.

Both JPEG and JPEG 2000 compressions essentially operate using a rather similar workflow. The fundamental differences between JPEG and JPEG 2000 are:

- To guarantee a linear behavior of processing time (DCT grows proportional to $n \cdot \log(n)$, where n means the total number of pixels), JPEG images are tiled into 8×8 pixel blocks prior to applying the image transform. This procedure is responsible for JPEG's typical blocking artifacts (**Color Plate 18, p. 594**) at high compression rates, owing to the Gibbs phenomenon causing the Fourier transform to fail to converge uniformly at discontinuities.[18] For JPEG 2000, in contrast, the original images are tiled only if explicitly required. Instead, tiling into 64×64 blocks is performed *after* the image transform.
- Instead of the Fourier-based DCT (Discrete Cosine Transform) as used in JPEG, the new JPEG 2000 compression makes use of the wavelet transform, which has been well known since the beginning of the nineties. At first glance, wavelet decompositions might not intuitively be identified as orthogonal transforms. The basic orthogonal character of the wavelet transform results from iteratively applying linear filters to the original image, resulting in a full decomposition which is termed a multi-resolution representation,[20] as illustrated in **Color Plate 20, p. 595**. Nevertheless, the wavelet coefficients correspond to the Fourier coefficients after an FFT or a DCT.
- The quantization strategy in JPEG is based on the fixed values of a uniform quantization table which has been tuned by empirical optimization. The quantization of the multi-resolution representation in JPEG 2000 is slightly more complex, because it depends on the position and the number of the respective recursion level. Furthermore, the quantization step size also depends on the dynamic range of the sub-bands.[7]

According to widespread opinion, the improvements in JPEG 2000 are closely associated with the wavelet transform. However, this belief seems questionable, as the comparison of **Color Plate 21, p. 596** with **Color Plate 22, p. 597** shows. In order to visually estimate the compression power of DCT and DWT, an image $f(x,y)$ has been transformed into the frequency domain using either DCT or DWT, resulting in $F(u,v)$. The coefficients $F(u,v)$ are subsequently sorted by value. Its reconstruction into the spatial domain is accomplished by considering only a predefined percentage ε of the major coefficients, yielding $f'(x,y)$. The remaining smaller coefficients are discarded, i.e., set to zero. Since DCT as well as DWT divides a signal into its energy fractions (Parseval's theorem), the resulting $f'(x,y)$ holds ε percent of the most significant information from $f(x,y)$. This simplified compression scenario benefits from entirely omitting the tiling prior to the image transform which causes the Gibbs phenomenon and thus the blocking artifacts. The calculation cost is increased by a factor of three, in theory.

The procedure described has been applied with different ε ranging from 100% (which results in the original image) down to 0.025% (where only 1/4000 of the original image information is reconstructed, i.e., 65 coefficients from an original 262,144 color values). **Color Plate 21, p. 596** shows the results of the 512×512 "Lena" image which has been compressed accordingly by means of a DCT transform. **Color Plate 22, p. 597** presents the respective images after a DWT-based compression using the biorthogonal 9/7 wavelet filters, the standard for lossy JPEG 2000 compression. Both DCT and DWT likewise indicate that discarding up to 99% of the original image information yields no major image distortions, and after discarding up to 99.9%, Lena is still easily identifiable.

Experiments of this type imply the conclusion that orthogonal decompositions of wavelet type do not necessarily provide an observable benefit with regard to compression quality compared with the discrete cosine transform. Instead, excessive artifacts in JPEG images at high compression rates occur not because of the cosine transform, but rather owing to the tiling into 8×8 pixel blocks and the resulting discontinuities of $f(x,y)$. Tiling an image before compression results in blocking artifacts in JPEG 2000 as well, as illustrated in **Color Plate 23, p. 598**, where a tiling size of 32×32 pixels prior to JPEG 2000 compression to 1:48 (this can be effected using the 'tiles' option in JJ2000[21]) had been selected for the image to the right, whereas the image to the left shows the untiled JPEG 2000 compression at the same rate. With reference to this experiment, unlimited block sizes in JPEG entailing the prevention of the Gibbs phenomenon would possibly allow compres-

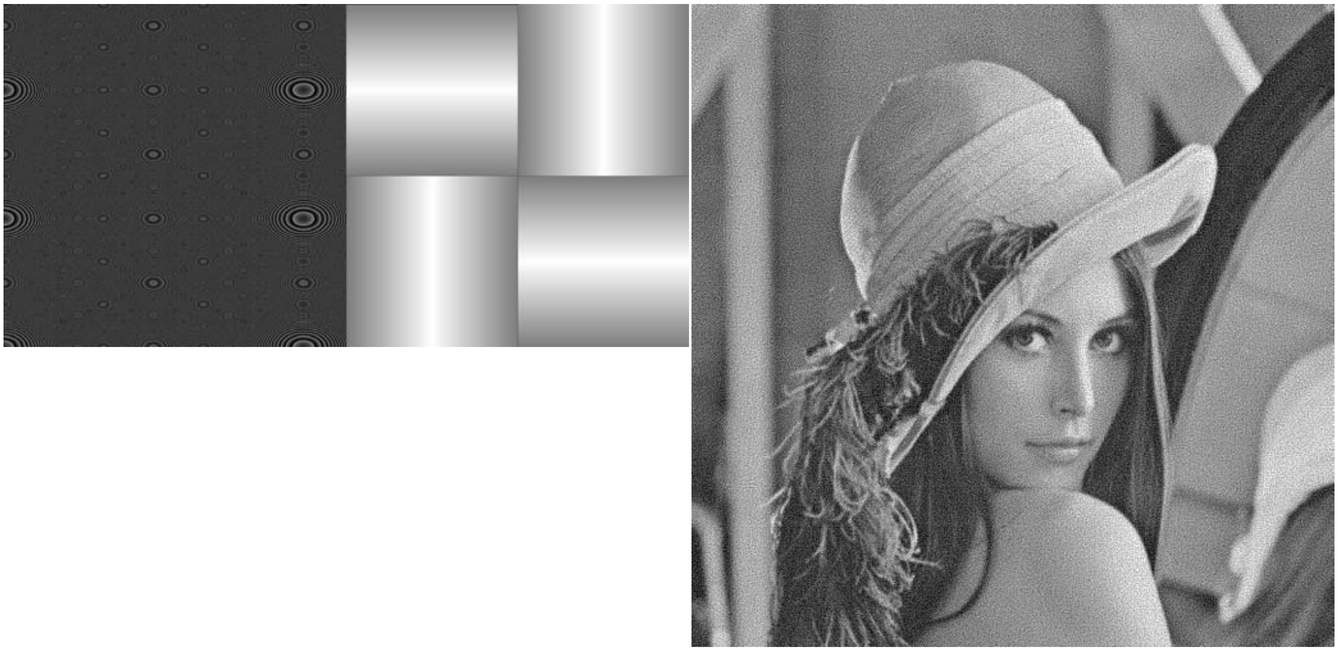

Figure 1. Test images "Unicolored" (left) and "Screened Lena" (right).

sion qualities comparable with JPEG 2000 even at high compression rates.

Perceptive JPEG and JPEG 2000 Rating

Different types of interactive visual ratings have been performed to achieve reliable quality estimations at varying compression intervals. A number of fundamental questions have been investigated, namely:

1. At which compression rates is it possible to recognize differences between the compressed image and the original?
2. Which compression rates of JPEG and JPEG 2000 result in a visually comparable reproduction quality?
3. At which compression rates does JPEG 2000 outperform JPEG and by what ratio?

The basis for our visual estimations are perceptive comparisons between the original image and its reproduction, similar to a proofing situation in the graphic arts industry. All tests were run under predefined viewing conditions and without time limits. Only high resolution images, typically of 1024×1280 pixels, have been considered.

In accordance with the mainly unattended use of image formats in common industrial workflow, the study focuses exclusively on the JPEG/JPEG 2000 core coding system[11] together with its default settings. We expect most users will not change those settings, since some of the optional parameters require more profound knowledge of the basic compression techniques. In our tests we are thus applying the JJ2000 4.1[21] software without any further options, except variation of the compression rate which is assigned with the *'rate'* option. It is, however, necessary to clarify that JPEG 2000 can be individually improved by going beyond the default parameters and by using visual frequency weighting, as mentioned in Annex J of Ref. 11.

Technical Environment

In order to achieve optimal reproduction accuracy, the highest quality IBM TFT-monitor IBM T221 was selected as a test platform. Its display area of 3849×2400 pixels at a display size of 22.2", i.e., 200 ppi, a contrast ratio of 400:1 and a brightness of up to 235 cd/m^2 guarantee natural viewing conditions comparable with high fidelity printing. The horizontal/vertical viewing angles are satisfactory, at 170°.

The software platform was an Intelli Station Pentium 4 running under Windows 2000. The test software had been written in JAVA. The applied JPEG compression software originates from the corresponding modules of JAI (Java Advanced Imaging). The JPEG 2000 images had been calculated with the help of JJ2000 4.1, the official standard reference software of the JPEG 2000 consortium.[21] All default settings had, however, been retained for image compression in both standards (see previous section). The only parameters altered were the compression quality in JPEG (which must be selected from the range of [0..1]) and the compression rate in bits per pixel in JPEG 2000 (which is preset in JJ2000 using the 'rate' option).

Test Images

More than one hundred different test images were inspected. The image size ranges from about 1 to 7 MB. Basically, all images had been chosen at a resolution allowing its pixels to be mapped exactly to one corresponding pixel on the TFT monitor. This is required

- to avoid additional filtering owing to pixel interpolation,
- nevertheless to take the best possible advantage of the hardware facilities,
- but to prevent further distortions due to the constrained monitor resolution.

Basically, two image collectives have been investigated, a small one containing about ten different images, examined using fairly accurate but time-consuming visual analysis, and a large one containing more than a hundred images; however, these were inspected in a less time-consuming way.

The small collective includes some of the ISO 12640 standard color image data downsized to 1024×1280 pix-

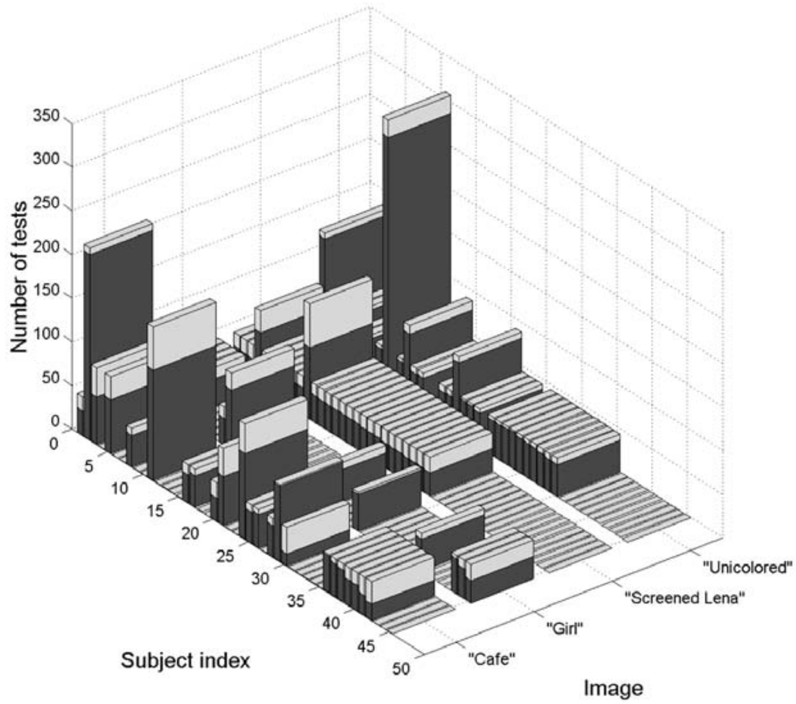

Figure 2. Overview of the completed test sets adapted to find the compression rates limits still yielding perceivable distortions. The examiners are specified by their examiner ID on the abscissa, the respective test image names on the ordinate. Each cuboid represents the total number of tests per image and examiner. The black lower parts indicate the number of correctly recognized tests while the gray parts denote the number of failures.

els each, in particular those images in **Color Plate 24, p. 598**, plus some artificial test images (Fig. 1) especially sensitive to distortions arising from JPEG and/or JPEG 2000 compression. This image collective forms the basis for all investigations in the next two sections.

The large image collective (136 images) evaluated in section 'JPEG/JPEG 2000 quality matching' is outlined in **Color Plate 25, p. 599**. It comprises:

1. all ISO-12640 images, which were also used as major test images when the JPEG 2000 standard was being drawn up,
2. some delicate artificial images as in the first collective, plus
3. many natural color pictures from various photo CDs, including
   - miscellaneous photographs from nature, landscapes and the arts,
   - some satellite images as well as cutouts from tourist maps, where even slight distortions are easily identifiable,
   - pictures containing areas of smooth color gradients, critical mainly for JPEG images,
   - shots of people, delicate in respect to defects in skin colors,
   - and pictures containing structural patterns.

Discrimination Test

The most evident requirement of a successful compression standard is that it should be capable of representing the original image without visually perceivable deficiencies. Obviously, this is possible only up to a certain compression limit. A discrimination test has been designed to investigate this issue.

The corresponding test software presents a set of three images on our IBM high-resolution monitor. The set comprises the original image, plus either its JPEG or its JPEG 2000 version. The supplementary third image is a duplicate of one of them. The three images are displayed in a random sequence next to each other (see layout in **Color Plate 26, p. 600**). The examiner's job is now to identify the unpaired image, knowing neither the rate nor the type of the applied compression. The result, which is either a correct or an incorrect answer, is recorded automatically together with the compression type and rate, plus identification of the examiner.

A group of 46 different testers performed a total of 141 test runs, each run comprising 22 confrontations of the original image with its JPEG version, as well as 22 confrontations with its JPEG 2000 version. Most of the testers went through one single test run only; some of them performed 2–7 iterations to improve data reliability. Each test run typically lasts about one hour and requires a high degree of concentration. The overall expenditure resulted in a total of 6204 comparisons. The test runs are illustrated in Fig. 2, which gives an overview of all acquired tests, grouped by examiners and test images. The height of each cuboid corresponds to the total number of test comparisons. The black section represents the portion of the correct answers, the gray section displays the number of incorrect recognitions. The compression rates considered had been randomly selected from a lookup table accommodated to each specific image, in order to concentrate on compression rates of sensitive magnitudes.

The examiner had to make a decision on one of the three images even if he was not able to perceive the distortions, these often being marginal. The probability of a correct identification is 1/3 in such cases, implying that only success rates considerably higher than 33% are statistically relevant.

A first impression of the results is given in Fig. 3. The diagram illustrates the ratio of correct identifica-

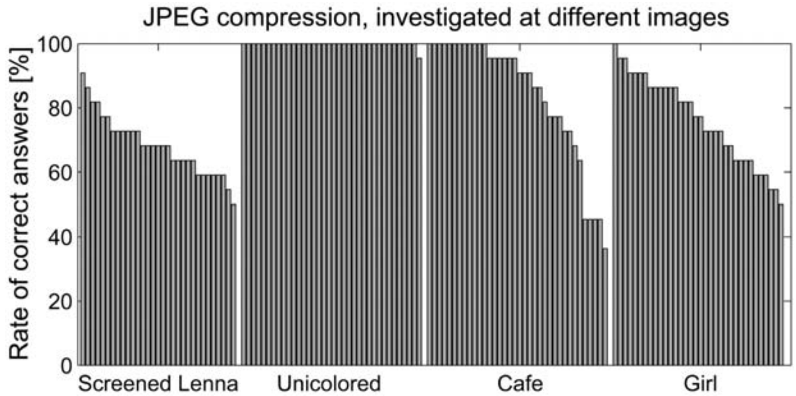

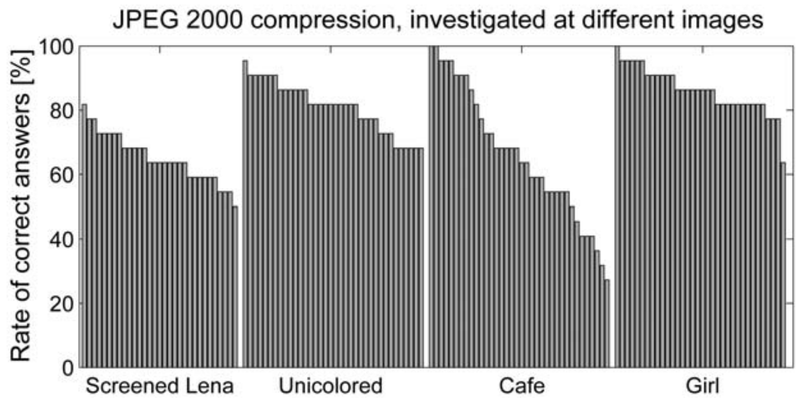

Figure 3. Success rates for all test sets grouped by images and compression technique (top: JPEG, bottom: JPEG 2000) and arranged in decreasing order of success quota. Each bar represents a test set containing a total of 22 recognition tests. Note the large scattering of the success rates, varying from 33 to 100 percent, implying large intra individual differences in power of visual discrimination.

tions per test image and test run. Each bar represents the percentile success rate of a test set containing 44 single recognition tests, each at different compression rates. The test results are sorted by success rates in decreasing order. JPEG and JPEG 2000 test sets are displayed separately. The inter-individual scattering is quite remarkable. While some testers had been able to accurately recognize up to 100% of the presented cases, others only accomplished a success quota close to the random success rate (33%) (It is debatable whether those subjects suffered from eye deficiencies or whether they did not essentially grasp the test procedure). Since the examined compression rates had been adapted to the specific image, the mean success quotas achieved (Table I, parenthesized values) are not expressive. Instead, the scattering values, σ, in Table I indicate how evidently the distortions could be perceived depending on the image.

The evaluation of the "Cafe" image (third group of bars) proved to be rather demanding, since its defects caused by both JPEG and JPEG 2000 compressions are rather subtle. This induces a large inter-individual standard deviation ($\sigma = 20\%$) of the identification rate. Conversely, the "Unicolored" artificial image responds more sensitively to both compression types, causing structural impairments on JPEG 2000 ($\sigma = 8.0\%$) and, above all, conspicuous color shifts on JPEG images, most easily recognizable even at the lowest compression rates ($\sigma = 0.8\%$). The "Girl" image results show a higher standard deviation on JPEG ($\sigma = 13.7\%$) than on JPEG 2000 ($\sigma = 7.2\%$) compression, meaning that JPEG 2000 is generating more distinct artifacts than JPEG for this specific image, astonishingly.

TABLE I. Standard deviations and mean values (parenthesized) of the success rates in percent, achieved on the JPEG and JPEG 2000 examinations of the four images investigated in Fig. 3.

Image	Std. deviation σ (mean) JPEG	Std. deviation σ (mean) JPEG 2000
Screened Lena	9.4 (68.3)	7.6 (65.1)
Unicolored	0.8 (99.9)	8.0 (80.9)
Cafe	19.5 (84.1)	20.4 (66.2)
Girl	13.7 (76.2)	7.2 (86.1)

An evaluation of the data presented in Fig. 3 is given in Fig. 4, showing the percentile quota of correct identifications, dependent on the compression rate (solid curves for JPEG, dashed curves for JPEG 2000). The results of all testers has been averaged for each compression rate, and the average values interpolated with smoothing splines.

As indicated in the previous section, the "Unicolored" artificial test image proves to be most sensitive especially for JPEG compression, allowing correct identification for nearly all compression rates and examiners. The "Screened Lena" and the "Cafe" images show correct identification of the distortions earlier in JPEG than in JPEG 2000, as could be expected, provided that JPEG

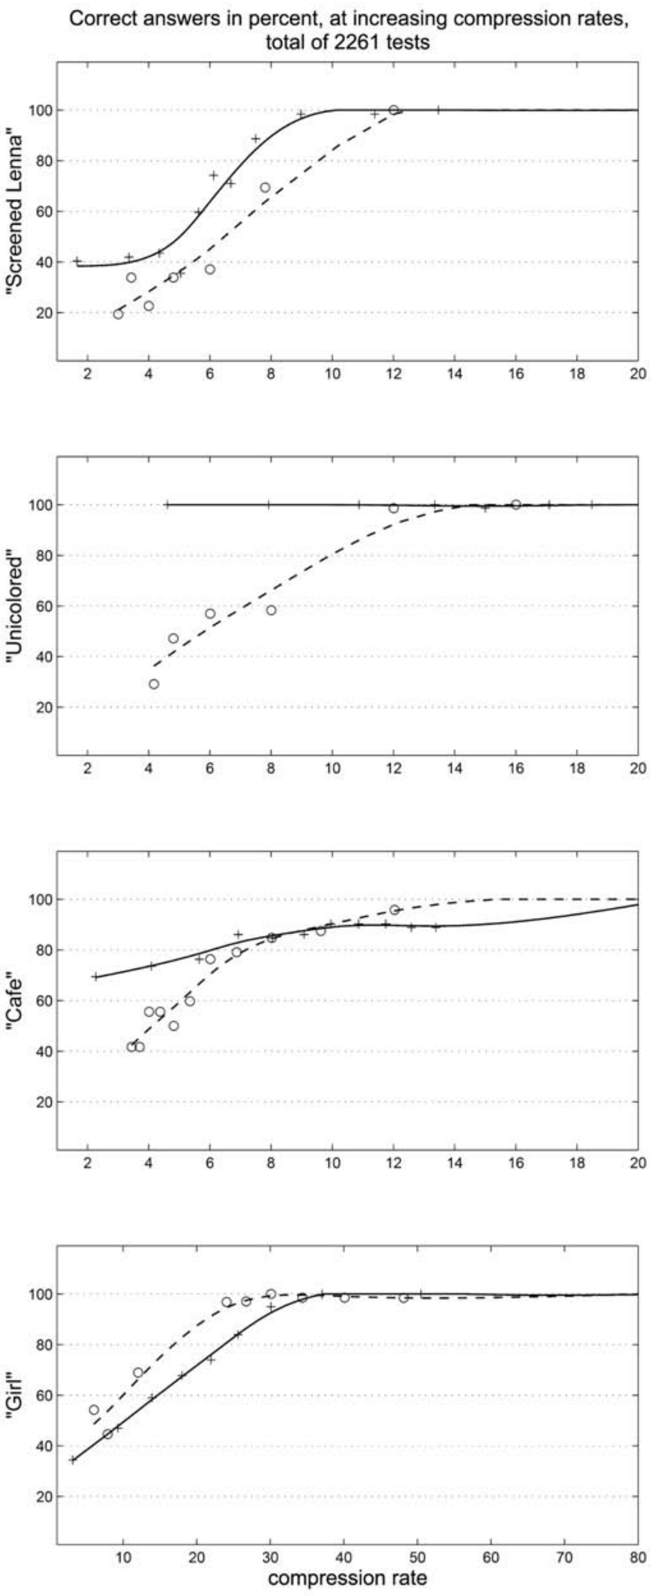

Figure 4. Percentage of correct identifications depending on the compression rate. The solid lines represent the results for the JPEG compression, the dashed ones for the JPEG 2000. Please note that a different scale is used for the "Girl" image.

TABLE II. Compression Rates at Specific Success Thresholds of 66%, 75% and 90% Success Quotas.

Image	Threshold	Rate α JPEG	Rate β JPEG 2000	Gain α [dB]
Screened Lena	66%	6.51	8.02	1.16
	75%	6.77	8.96	1.22
	90%	8.04	10.69	1.24
Cafe	66%	0.97	5.59	7.59
	75%	4.54	6.50	1.56
	90%	14.82	9.88	- 1.76
Girl	66%	17.38	11.94	- 1.63
	75%	21.48	15.01	- 1.56
	90%	28.56	21.25	- 1.28
Mean	66%	7.66	8.50	0.46
	75%	10.21	10.24	0.01
	90%	15.00	14.23	- 0.51

2000 is in fact improved in quality compared to JPEG. The "Girl" image, however, produces the contrary message, indicating that its JPEG distortions are even less conspicuous than its JPEG 2000 artifacts at low compression rates.

Table II shows the compression rates (α for JPEG, β for JPEG 2000) at specific success thresholds of 66% (twice the random success rate), 75% and 90% (100% was not always achieved due to lack of assiduity). It is striking that the differences between JPEG and JPEG 2000 seem to be smaller than the differences between the images. Additionally, the gain, ρ, in decibels

$$\rho = 10 \cdot \log_{10}\left(\frac{\beta}{\alpha}\right) \tag{1}$$

of the JPEG 2000 compression rate β compared to the JPEG compression rate α is listed in Table II. Negative gains, signifying that JPEG is superior to JPEG 2000, appear surprisingly often.

Figure 5 illustrates the quota of all examiners providing correct answers up to a specific compression quality. In other words, the criterion applied was the compression rate at which the examiner was for the first time failing to recognize the correct image. This analysis basically yields similar results to Fig. 4.

To summize,

1. JPEG 2000 images do not show any significant benefit compared with JPEG images with respect to its perceivable distortions at low compression rates.
2. The comparative situation is very image dependent.
3. In some cases, JPEG might even outperform JPEG 2000 in terms of high quality compression power.

JPEG/JPEG 2000 Comparison Test

Medium and high compression rates are the major purpose of lossy image compression. Such compressions necessarily yield visually perceivable distortions. Estimating the compression quality sequentially requires perceptually weighting its image defects. Though this is a subjective task, visual rating of images is part of the day-to-day business of the graphic arts industry, and our examiners are familiar with this issue. As is customary, they presumed a reproduction was optimal if it perceptually appeared to be as close to the original as possible.

Correspondingly, our second test addresses the question of visually comparable JPEG and JPEG 2000 re-

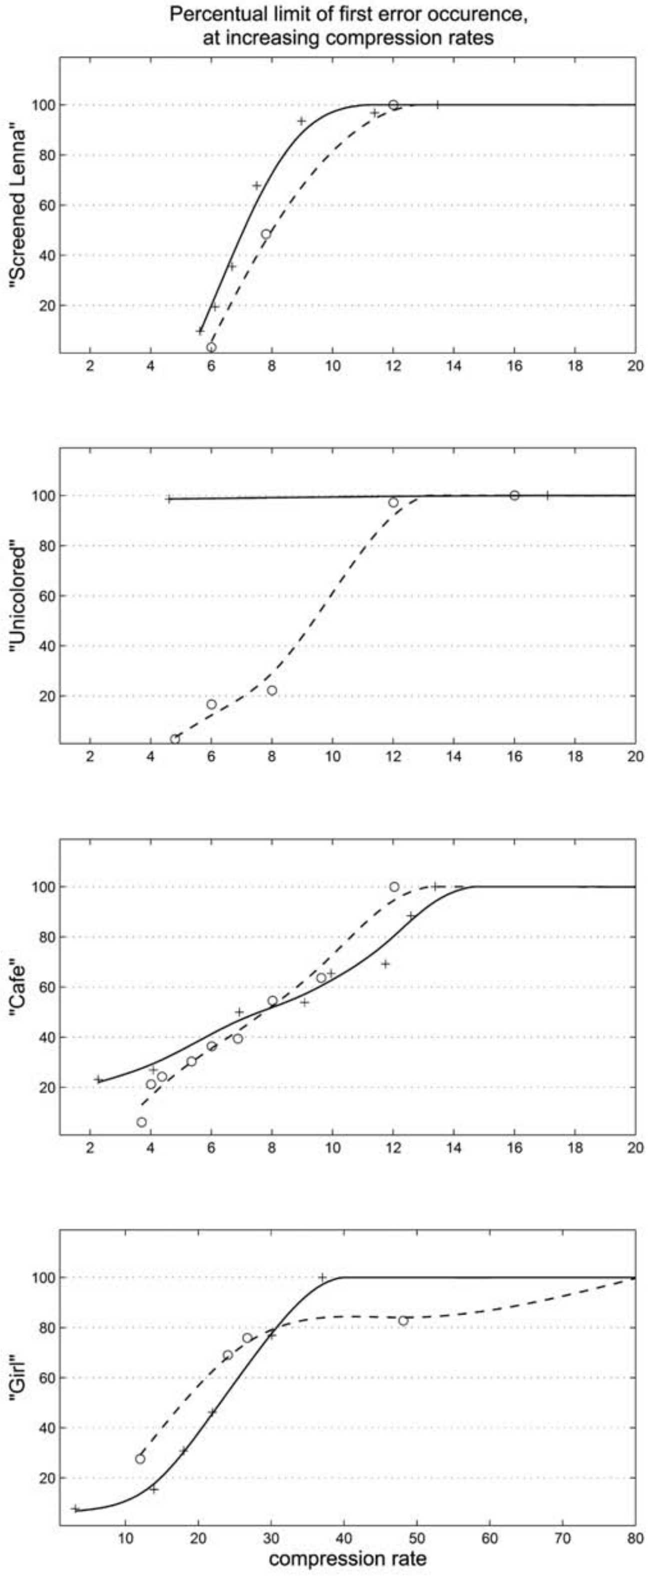

Figure 5. Quota of the testers in percent having successfully identified the differing images up to a specific compression rate, i.e., their first differentiation error occurred at the next lower compression rate. The solid lines represent the results for the JPEG compression, the dashed ones for the JPEG 2000. Please note that a different scale is used for the "Girl" image.

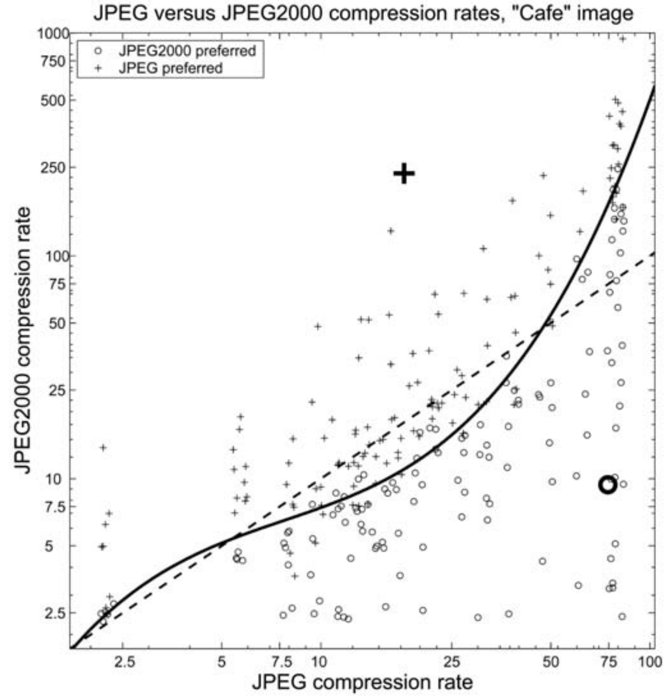

Figure 6. JPEG and JPEG 2000 images at different compression rates are compared by an examiner with regard to reproduction quality. Crosses stand for the cases where the examiner favored the JPEG version, circles mean that he preferred the JPEG 2000 version. The test image examined is the ISO image "Cafe" (**Color Plate 24, p. 598**, center). The basic values of the preference borderline are interpolated by a polynomial fit of order 3. The dashed line denotes the location of equivalent JPEG/JPEG 2000 compression rates. The bold cross and circle represent the specific examples given in **Color Plate 28, p. 601.**

production qualities. Let α (β) be the JPEG (JPEG 2000) compression rates applied to an image, resulting in a perceptive reproduction quality $Q(\beta)$ ($Q(\alpha)$). Consequently, for each pair (α,β) we obtain

$$Q(\alpha) \leq Q(\beta) \text{ or } Q(\alpha) > Q(\beta) \qquad (2)$$

denoting the visual preference of JPEG 2000 or of JPEG respectively. The conditions (2) divide the plane spanned by α and β into two sectors, where either one of the standards predominates. The line dividing the two sectors — we call it the *preference borderline* — can be understood as the location of equivalent reproduction qualities $Q(\beta) \approx Q(\alpha)$.

The test software configuration is displayed in **Color Plate 27**, **p. 600,** showing a JPEG together with a JPEG 2000 compressed image at randomly varied compression rates in the range from 1:1 to 1:1000. They are randomly sequenced in order to hide the compression type to the examiner, as far as possible. Both images might be characterized by minor to serious distortions. The examiner is required to choose his preferred image, in terms of reproduction quality. This is quite a sophisticated task, especially at high compression rates, since JPEG 2000 defects are of a substantially different character from JPEG distortions. Hence, only examiners actually trained in judging image damage (photographers, experts in the printing and graphic arts industry, image processing professionals) were selected to perform this kind of rating.

In order to speed up the evaluation process and to avoid obvious quality contrasts, the test software includes a self-learning strategy: the program subsequently constricts the compression rate bandwidths of the presented JPEG/JPEG 2000 pair. Moreover, it controls the compression rates in question in such a way that a balanced number of JPEG relative to JPEG 2000 preferences can be acquired. For instance, if the number of JPEG preferences is too low, the program automatically responds by increasing the JPEG 2000 compression rates in question, in order to give JPEG a better chance.

Each examiner was instructed to keep on running the test program until enough data points had been achieved to clearly separate JPEG 2000 from JPEG preferences. Generally, a total of 200–500 decisions was made per examiner and image. This typically required about 3–6 hours of concentrated work. The test images displayed in **Color Plate 24, p. 598**, plus the two images in **Color Plates 18 and 19, p. 594**, which are especially delicate for JPEG and JPEG 2000 distortions, were each evaluated by up to 5 different examiners .

Figure 6 shows an evaluation example of the "Cafe" image (**Color Plate 24, p. 598**, center), provided by one single examiner. The JPEG compression rate is mapped to the abscissa, the opposed JPEG 2000 compression rate to the ordinate. Each dot stands for a single decision. Crosses represent a preference for the JPEG, circles for the JPEG 2000 image.

To better illustrate the meaning of this diagram, the bold cross and circle exemplify two JPEG/JPEG 2000 compression pairs which are explicitly highlighted in Fig. 6. The top image pair of **Color Plate 28, p. 601** represents the "Cafe" image compressed at the rates emphasized with the bold cross. It illustrates an unequivocal case where the quality of the right hand JPEG image outmatches the quality of the left hand JPEG 2000 image. Conversely, the bold circle in Fig. 6 denotes the compression rates of the two lower images of **Color Plate 28, p. 601** where the JPEG 2000 compression on the left outperforms JPEG on the right. The compression rates of these two examples are quite far apart from each other, in order to ensure that the quality gradients become distinctly visible at medium printing quality and at the far too small reproduction size. On the highest resolution IBM T221 monitor, the quality differences are definitely much more distinct.

Crosses and circles slightly overlap because of contradictory statements due to the intuitive nature of visual decisions. Standard statistical methods such as logistic regression proved to be unsuitable for calculation of the *preference borderline*, because on one hand, the $\alpha, \beta,$ tuples are not normally distributed, and on the other hand no standard function for the borderline characteristic, which is most image dependent, could be presumed.

The preference borderline function was determined from the raw data by applying a two-stage process. First, for a given JPEG compression rate α, a transition point $t(\alpha)$ separating JPEG from JPEG 2000 preference is determined. If the JPEG and the JPEG 2000 preferences did not overlap, $t(\alpha)$ was set to the epicenter of the flanking regions separating crosses from circles. If overlapping occurred, the barycentric balance point within the overlapping interval was taken. In a second step, the samples calculated in this way were interconnected using a least squares polynomial fit, yielding the preference borderline.

Figure 7 shows the results of all the evaluated images ("Girl", "Cafe" and "Champagne" in **Color Plate 24, p. 598**, "Frog" and "Sand" in **Color Plates 18 and 19, p. 594**). The dotted preference borderlines represent the evaluations per tester. Even though the expertise of all the examiners' in judging reproduction quality was respectable, the inter-individual scattering of the compression rates by factors of up to 4–8 is quite remarkable. The overall average preference borderlines denoting equivalent quality of JPEG and JPEG 2000 are displayed in solid bold. The total numbers of collected judgments, denoting either JPEG or JPEG 2000 preference, amount to 754/798 for the image "Girl", 560/666 for "Cafe", 437/697 for "Champagne", 154/218 for "Frog" and 288/278 for "Sand".

The dashed line, illustrating equivalent JPEG and JPEG 2000 compression rates, divides the (α, β) plane into two moieties, the upper one denoting JPEG 2000 compression greater than JPEG compression, and the lower one vice versa. This implies that all parts of the preference borderline located above the dashed line stand for compression bands where JPEG 2000 outperforms JPEG, whereas parts below represent bands with JPEG predominance.

All the evaluations in Fig. 7 indicate an indisputable quality benefit of JPEG 2000 at high compression rates above 50. Hence, JPEG 2000 clearly proves to be more suitable than JPEG for very low quality requirements, such as for internet applications. It must, however, be emphasized that images compressed at the highest rates are of rather poor quality, owing to the dramatic information loss.

Medium compression rates clearly below ~ 40 occasionally show an utterly unexpected trend where JPEG compression outperforms JPEG 2000 (images: "Sand", "Girl", "Cafe", "Champagne"). The examiners' justification for this rather surprising fact has been explained on one hand by troublesome cloudy distortions reminiscent of blurring (see **Color Plate 29, p. 602**) whereas JPEG distortions in contrast show a rather grainy character suggesting sharpness. On the other hand, JPEG 2000 tends to lose structure in fine-grained local image areas (see the "Sand" image in **Color Plate 19, p. 594**). We shall discuss this issue in the following sections.

Low compression rates of less than 3 often do not include any visible distortions. Thus the preference borderline obtained at lowest compressions is probably not representative. This assertion is, however, highly image dependent.

Figure 8 shows the relative gain ρ of JPEG 2000 over JPEG compression rates in decibels, for the images "Girl", "Cafe", "Champagne", "Frog" and "Sand". The characteristics were calculated from the average preference borderlines of all testers (bold solid lines in Fig. 7), using Eq. (1). The "Frog" characteristic (dotted) clearly indicates that JPEG 2000 provides performance superior to JPEG, which causes heavy tiling artifacts, except for lowest compression rates (see **Color Plate 18, p. 594**). However, all the remaining images feature distinct bands at medium compression rates, where JPEG outperforms JPEG 2000. This is particularly conspicuous on the "Sand" image (dashed characteristic), showing negative gains of more than -3 dB due to structural loss. At high compression rates, however, unaesthetic color defects with JPEG overtake its JPEG 2000 artifacts (see **Color Plate 19, p. 594**).

In brief,

1. JPEG 2000 clearly outperforms JPEG at high compression rates above 50.

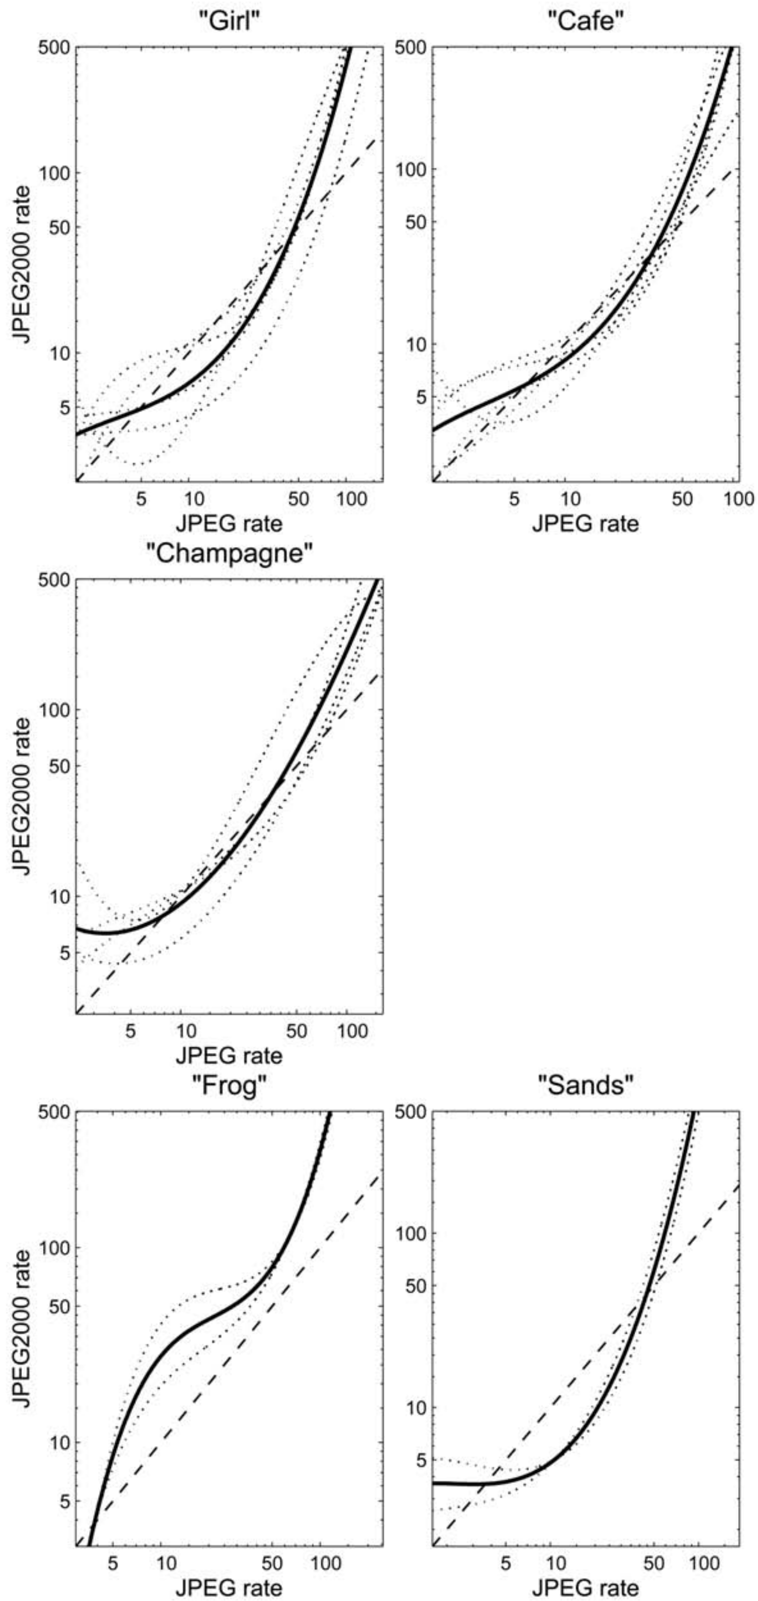

Figure 7. Multiple JPEG to JPEG 2000 preference borderlines (dotted) of different images and examiners. The bold solid borderlines represent the average overall examinations of the specific image. The dashed line shows the location of equivalent compression rates for JPEG and JPEG 2000.

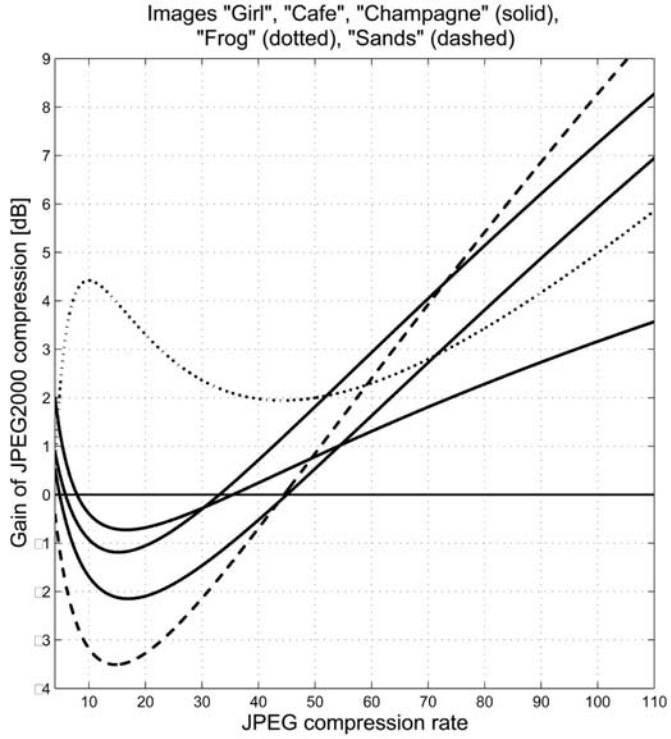

Figure 8. Gain in dB of JPEG 2000 over JPEG compression, for the images "Girl", "Cafe", "Champagne", "Frog" and "Sand", examined by multiple testers. While JPEG 2000 undoubtedly results in better qualities than JPEG for the "Frog" image (dotted line), JPEG distinctly outperforms JPEG 2000 at medium compression rates for all remaining images.

2. At medium rates, however, lots of the examples considered surprisingly produce precisely the opposite result.
3. The comparative situation is heavily dependent on the image.

JPEG/JPEG 2000 Quality Matching

The examinations in the previous two sections unexpectedly imply the hypothesis that JPEG could be superior to JPEG 2000 at medium compression rates below ~ 40, in terms of perceptive qualities. Owing to the large scatter between different images, in conjunction with the small number of evaluated images, statistical relevance of the latter speculations is, however, absent.

Hence, in order to confirm the suspicion thus aroused, a more straightforward and less time consuming type of evaluation, optimized for analyzing large image collectives, has been designed. The intention of the new evaluation procedure is to gather preference borderline characteristics of images, similar to the results in section 'JPEG/JPEG 2000 comparison test', but at only a fraction of the labor costs.

For logical reasons, the preference borderlines should not be considered until the lowest compression rates at which the examiner first perceives any distortions. Its ending point is represented by the highest possible JPEG compression rate, since the JPEG 2000 reference software[21] allows compressions up to a quality loss far beyond the lowest JPEG compression quality (usually already terribly low). In between these start and end points, the characteristics are recorded by the test program with a reasonable resolution.

To speed up the task of obtaining a preference borderline as efficiently as possible, a rather pragmatic mechanism has been developed to find the JPEG compression rate opposing a predefined JPEG 2000 compression. Given a JPEG 2000 reference image at a certain compression rate, the quality of the opponent JPEG image is interactively varied by the examiner by adjusting a slider which increases or reduces the JPEG compression rate, insofar as its quality begins to match those of the tag image.

More specifically, the evaluation workflow operates as follows:

1. For both JPEG and JPEG 2000, the compression rates yielding reproductions barely but nonetheless still distinguishable from the original image are investigated. Those limits determine the starting point of the preference borderline.
2. Considering the JPEG compression at its highest compression rate, the JPEG 2000 opponent of a comparable quality is estimated. This task, featuring the preference borderline end point, is rather delicate, inasmuch as any distortions caused by JPEG are of an entirely different character from those caused by JPEG 2000.
3. The compression rate interval between the preference borderline's start and end points is sliced into 10 equally spaced interim compression values. A logarithmic scale was used, yielding the best linear perceptual quality progression.
4. For each interim compression rate, the JPEG/JPEG 2000 tuples, coinciding as closely as possible in terms of reproduction quality, are subsequently determined.

This procedure typically makes it possible to determine the preference borderline of a single image within 10–15 minutes. Thus, the large image collective presented in the section 'Test Images' could be evaluated by different examiners within a reasonable time frame.

Figure 9 shows the preference borderlines of all 136 images in the large image collective (see section 'Test Images'), evaluated by a single examiner. Each characteristic is defined by a start and an end point, including 10 interim points. They represent the location of equivalent distortions due to JPEG (abscissa) and JPEG 2000 (ordinate) compressions at varying compression rates. As in section 'JPEG/JPEG 2000 comparison test', the dashed black line denotes the position of equivalent compression rates of JPEG and JPEG 2000, implying that preference borderlines or parts of them located above this balance limit mean JPEG 2000 compression is preferable to JPEG and vice versa. The large bandwidth of the results in Fig. 9, associated with the variability of the image motifs, is quite remarkable. A median rather than a mean characteristic of all evaluations — the preference borderlines cannot be assumed to have a Gaussian distribution — is shown in bold black. In accordance with the observations in section 'JPEG/JPEG 2000 comparison test', the median characteristic crosses the balance limit at the compression rate of 51 making it evident that JPEG 2000 outperforms JPEG only *above* this compression limit, according to this specific examiner and focusing on our large image collective.

Preference Borderline Characteristics

For all the evaluations in Fig. 9, Fig. 10 presents the gain in decibels of JPEG 2000 over JPEG, with regard to the compression performance at equivalent quality (see

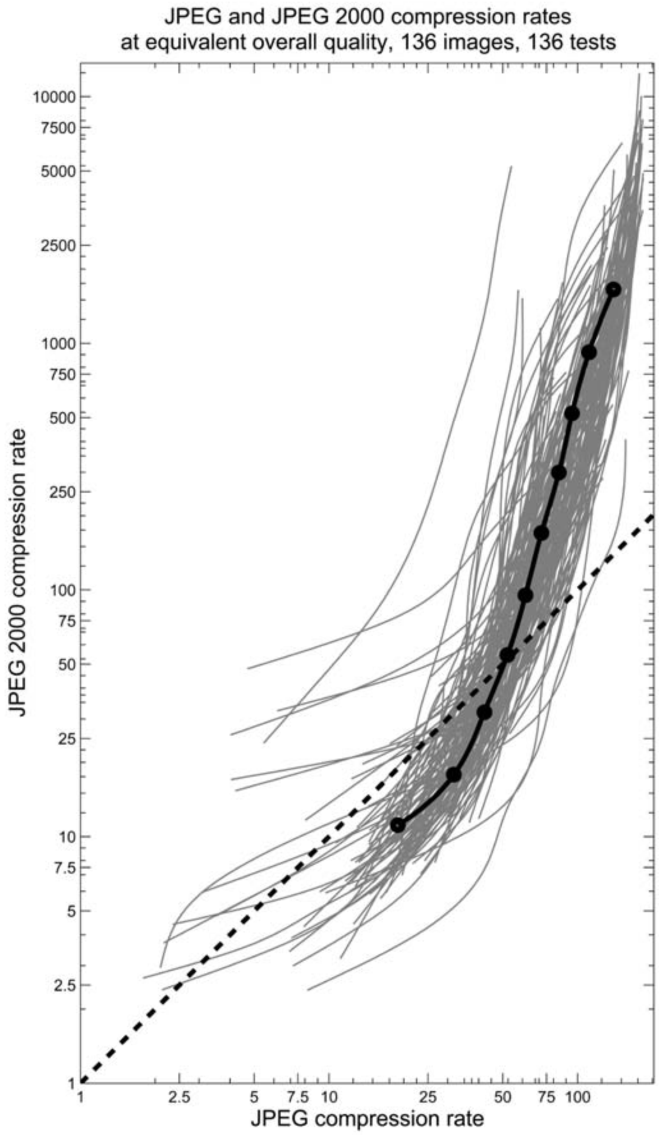

Figure 9. Preference borderline characteristics of 136 different test images (gray curves), determined by a single examiner. The median characteristic is plotted in bold black. For clarification, the line denoting equivalent compression rates is plotted in dashed black.

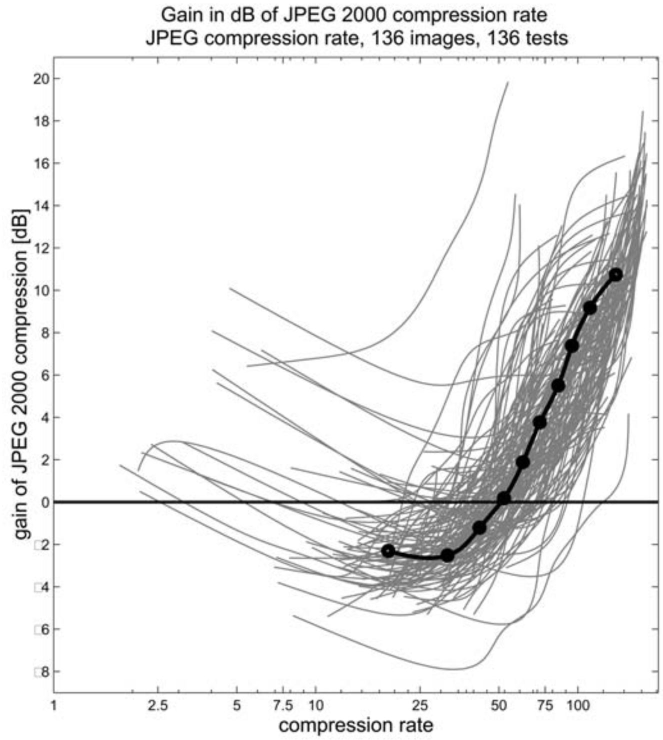

Figure 10. Quality characteristics determined by the same examiner and for the same 136 test images as in Fig. 9. The figure shows the gain in dB of JPEG 2000 in compression performance relative to the old JPEG. The median characteristic is plotted in bold black.

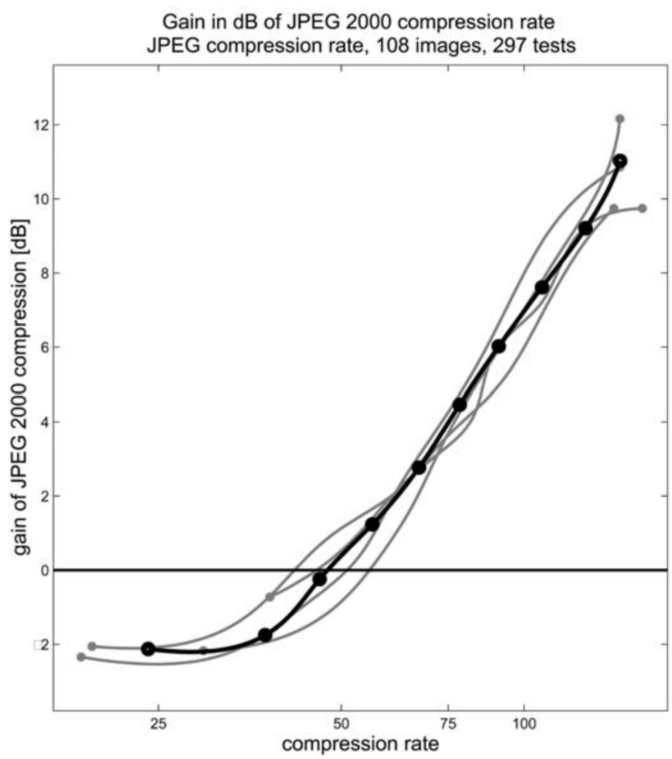

Figure 11. Median courses of the quality characteristics in dB determined by four different examiners (gray curves). The median characteristic over the grand total of 297 examinations of 108 different images is shown in bold black.

definition of gain in section 'Discrimination test'). Positive values represent the predominance of JPEG 2000, negative values the predominance of JPEG. The gain extends over a broad interval of (–8 dB, +20 dB). The median characteristic (plotted in bold black) indicates JPEG 2000 outperforming JPEG by up to 10 dB at the highest compression rates, meanwhile underperforming against JPEG by up to 2 dB at compression rates below 50.

A total of four examiners independently examined the large image set, or at least parts of it. The inter-examiner comparison is shown in Fig. 11. For each examiner, the median characteristic of his examinations is outlined. The bold black characteristic represents the median of all evaluations. The first fact which becomes obvious is the low variance among the examiners, ranging from $\rho = 0.74$ dB at low compression rates up to $\rho = 1.15$ dB at high rates; on average $\rho = 0.88$ dB. Consequently, the overall median characteristic is comparable with the

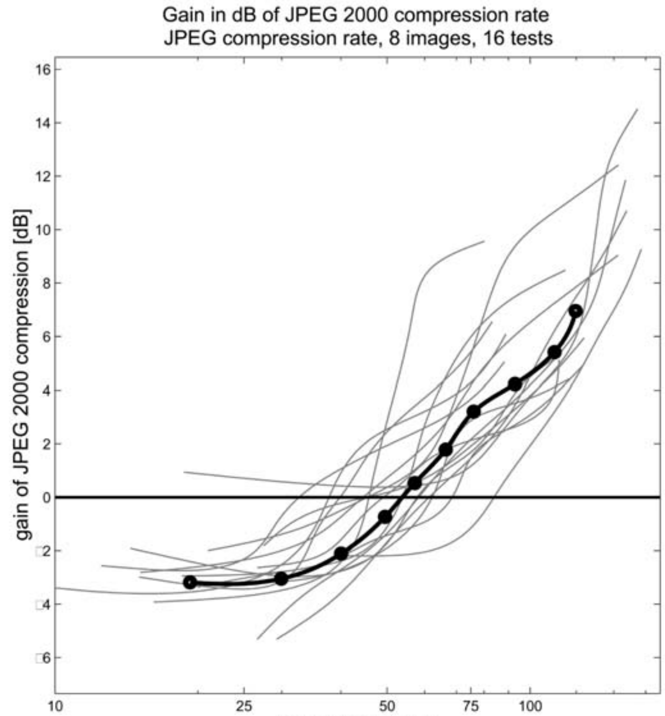

Figure 12. Quality characteristics of the eight ISO 12640 images, determined by a single examiner, plus its median characteristic (bold black).

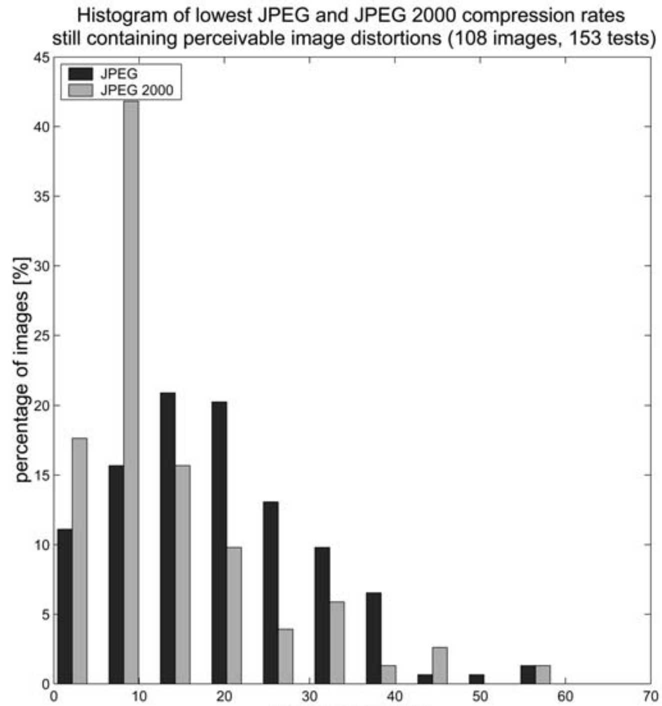

Figure 13. The lowest JPEG (black) and JPEG 2000 (gray) compression rate tuple of each preference borderline represents the limit where the examiner was not yet able to visually perceive any image distortions due to compression, i.e., the lowest discernible rates. The figure shows the relative distribution of these values for all test images.

evaluation in Fig. 10, favoring JPEG at low and JPEG 2000 at high compression rates, and crossing the balance limit of 0 dB at a compression rate of 47.4. A second important observation is that lower compression limits are rather dependent on the examiner. Two examiners (A and B) achieve a median lowest rate of 18.6 and 19.4 respectively, where they still notice deformations, whereas two examiners (C and D) end up at 29.6 and even 38.1.

The large image collection also includes the well known ISO 12640 standard, comprising eight images so far; these are of special interest, since they played a decisive role in the development process of JPEG 2000. They were originally in CMYK and have been transformed to RGB using Adobe Photoshop™. The image resolution of 1920 × 1536 pixels is just small enough to enable mapping of two images next to each other on the IBM T221 monitor for purposes of comparison. Their quality characteristics, evaluated by the same examiner as in Figs. 9 and 10, are displayed in Fig. 12. Qualitatively, the examinations yield the same conclusions as the results of the entire collective in Fig. 10: JPEG 2000 is favorable for compressions above 53 and is disadvantageous for rates below this. The trend is even more pronounced for the ISO images than for the entire collective, i.e., the median benefit for JPEG 2000 is only 8 dB at most, while the gain of JPEG increases to more than 3 dB at low rates.

Compression Limit of Perceivable Distortions

The first step in exploring the preference borderlines consists in identifying the lowest compression rates still producing visible artifacts. These lowest discernible rates might be compared between JPEG and JPEG 2000 in the same way as in section 'Discrimination test'. Though their statistical accuracy is evidently inferior to the latter ratings, they had been inspected on a large amount of images.

Figure 13 shows the histogram of the lowest JPEG (black) and JPEG 2000 (gray) compression rates for the large image collective, having been evaluated by examiners A and B, i.e., those two examiners having judged with low standard deviation. Obviously, the JPEG 2000 distribution already reaches its peak at lower compression rates than JPEG. Accordingly, the mean is 11.6 for JPEG 2000 and 19.23 for JPEG. The standard deviation is $\sigma = 11.4$ in either case. This implies that JPEG 2000 is showing small distortions even earlier than JPEG. Considering all four examiners, the lowest discernible rates average out at 14.4 for JPEG 2000 and 24.1 for JPEG, showing the same trend. The standard deviations are 16.4 and 14.5, respectively. To summarize, the histogram in Fig. 13 yields the same conclusions as the study in section 'Discrimination test'.

PSNR Considerations

Instead of performing visual studies, JPEG to JPEG 2000 comparisons have often been achieved so far by calculating the peak signal to noise ratio (*PSNR*), this being defined as

$$PSNR = -20 \cdot \log_{10}\left(\frac{RMSE}{2^b - 1}\right) \qquad (3)$$

where *RMSE* is the root mean square error and b the bit depth of the original image. This measure is often

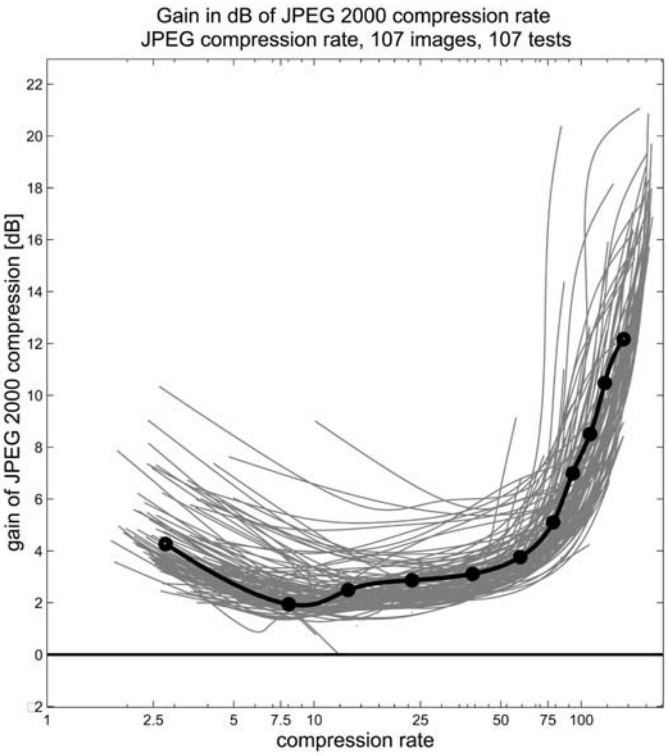

Figure 14. Locations of equivalent PSNR values for JPEG and JPEG 2000 compression for 108 different test images. The bold black characteristic represents the median characteristic.

chosen for assessing perceptual image distortions.[15] In order to compare our visual evaluations with the *PSNR* quality metric, the compression rates obtained by opposing a JPEG and a JPEG 2000 version at coincident *PSNR*s had been calculated, in contrast to subjective quality ratings. Figure 14 shows the characteristics emerging with *PSNR* estimations. In contrast to the preference borderlines arising from perceptive quality assessments (Figs. 10–12), *PSNR* considerations classify JPEG 2000 as clearly superior to JPEG for all images and compression intervals.

The basic question of the rather controversial results of the perceptual and the *PSNR* quality assessments can be explained by the specialty of Wavelet transforms to preserve spatial information. Due to this characteristic, the information loss effected by the compression is kept locally instead of being spread over the full image as in the case of the DCT transform of JPEG. Thus, distortions caused by JPEG 2000, especially those at low compression rates, typically appear notably localized. Thus, even though reproducing the major image parts in an excellent way, there are small regions containing subtle structures which are often neglected; nevertheless they are rather eye-catching when perceptually evaluated.

Conclusions

The goal of this work was to investigate the potential of JPEG 2000 in terms of reproduction quality and to compare it with the JPEG standard. The applied ratings were based on visual estimations on a highest-quality monitor. The evaluations achieved included

1. detecting the lower compression limits of visually perceivable deficiencies,
2. ranking JPEG and JPEG 2000 compression qualities relative to each other, yielding preference borderlines running from lowest to highest compression rates, with both tests at the best possible accuracy, but for a low number of images only. To overcome this drawback, a last evaluation was aimed at
3. gathering preference borderlines in a more time-saving and pragmatic way, allowing evaluation of more than one hundred images per examiner, including the ISO 12640 standard images.

The first two evaluations were based on a small number of images and yielded the following conclusions:

1. In terms of perceivable distortions, JPEG 2000 did not significantly outperform JPEG at low compression rates. Instead, the inverse case was occurring.
2. JPEG 2000 is capable of compressing at rates far beyond the scope of JPEG. The perceptive quality comparisons, moreover, acknowledge JPEG 2000 as being clearly superior to JPEG at high compression rates above 50.
3. Most unexpectedly, the quality of JPEG 2000 frequently underperforms against JPEG at compression rates below 50.
4. Those trends are highly dependent on the image motif, however. At medium compression rates, the inter-image scatter turned out to be considerably higher than the differences between the two compression standards.

The latter conclusion entailed a further type of evaluation, adapted for a large bundle of images. The hypothesis stated with evaluations 1 and 2 could be well corroborated by this third examination. Specifically, a compression rate of around 50 again proved to be a critical lower limit for JPEG 2000 predominance.

As expected, both standards, JPEG and JPEG 2000, feature major drawbacks at medium and high compression rates.

Characteristic types of JPEG distortions are:

1. Tiling at high compression rates caused by the fragmentation into 8 × 8 pixel blocks (**Color Plate 18, p. 594**)
2. Displeasing color shifts at high compression rates (**Color Plate 19, p. 594**)

JPEG 2000-specific deformations typically include:

1. Cloudy distortions at high compression rates reminiscent of blurring ("wavelet clouds", in **Color Plate 29, p. 602**)
2. Local texture drop-offs occurring in areas containing subtle structures (**Color Plate 30, p. 603**)

Summing up, JPEG as well as lossy JPEG 2000 compression techniques are questionable for high quality imaging, but they are well suited for medium and low quality applications, where JPEG 2000 especially has great potential. Owing to the highly motif-dependent scattering, the resulting reproduction quality is, however, rather unpredictable.

Beyond doubt, our fundamental statement is an unexpected affection for JPEG compression at medium rates below 50. This is the main application range for photography and the graphic arts industry, and is thus of major importance. Since JPEG 2000 artifacts are rather local at low compression rates, our results are contradictory to the generally accepted PSNR considerations for image quality assessment.

All in all, we believe that our extensive study has been carried out carefully and evenhandedly. We encourage readers to verify our findings, which are of course subjective, and would welcome any responses.

Acknowledgment. This work has been partly supported by cooperation from Ifra (International Newspaper Color Association) in Darmstadt, Germany, and by Ugra (Association for the Promotion of Research in the Graphic Arts Industry) c/o EMPA in St.Gallen, Switzerland. Sincere thanks to Brigitte Bänziger, Urs Bünter, Karl Kurz, Walter Steiger, and to many other coworkers, for their highly qualified evaluations and for spending numberless hours examining the extensive image material. Sincere thanks to Edward D. Anwyl, Display Business Unit IBM, for his support for the project.

References

1. ISO/IEC 10918-1, *Information Technology - Digital Compression and Coding of Continuous-tone Still Images - Part 1: Requirements and guidelines*, 1994.
2. ISO/IEC 10918-2, *Information Technology - Digital Compression and Coding of Continuous-tone Still Images - Part 2: Compliance Testing*, 1995.
3. W. Pennebaker and J. Mitchell, *JPEG Still Image Data Compression Standard*, Chapman and Hall, London, UK, 1993.
4. G. K. Wallace, The JPEG Still Picture Compression Standard, *IEEE Trans. Consumer Electronics* **38**(1), 18–34 (1992).
5. ISO/IEC JTC 1/SC 29/WG 1 FCD 14495 (public draft), *Lossless and near-lossless coding of continuous tone still images (JPEG-LS)*, July 1997.
6. M. Charrier, D. Santa Cruz and M. Larsson, JPEG2000, the Next Millenium Compression Standard for Still Images. In *Proc. of the IEEE International Conference on Multimedia Computing and Systems (ICMCS),* vol. 1, IEEE Computer Society, Los Alamitos, CA, 1999, pp. 131–132.
7. C. Christopoulos and T. Ebrahimi, The JPEG2000 Still Image Coding System: An Overview, *IEEE Trans. Consumer Electronics* **46**(4), 1103–1127 (2000).
8. M. W. Marcellin, M. Gormish, A. Bilgin, and M. Boliek, An Overview of JPEG 2000, in *Proc. IEEE Data Compression Conference (DCC'2000)*, IEEE Computer Society, Los Alamitos, CA, 2000, pp. 523–541.
9. A. N. Skodras, C. Christopoulos and T. Ebrahimi, JPEG2000: The Upcoming Still Image Compression Standard, in *Proc. 11th Portuguese Conf. Pattern Recognition (RECPAD 2000)*, Elsevier/North-Holland, Amsterdam, 2000, pp. 359–366.
10. D. S. Taubman and M. W. Marcelling, *JPEG2000 Image Compression Fundamentals, Standards and Practice*, Kluwer Academic Publishers, Boston/Dordrecht/London, 2002.
11. ISO/IEC 15444-1, *Information technology - JPEG2000 image coding system - Part 1: Core coding system*, http://www.jpeg.org/JPEG2000.html, 2000.
12. D. Santa Cruz and T. Ebrahimi, An analytical study of JPEG 2000 functionalities, in *Proceedings of the IEEE International Conference on Image Processing (ICIP)*, vol. 2, IEEE Operations Center, Piscataway, N.J., 2000, pp. 49–52.
13. I. Moccagata, S. Sodagar, J. Liang, and H. Chen, Error Resilient Coding in JPEG-2000 and MPEG-4, *IEEE (JSAC)* **18**(6), 899–914 (2000).
14. D. Santa Cruz, T. Ebrahimi, M. Larsson, J. Askelöf, and C. Cristopoulos, Region of Interest Coding in JPEG2000 for interactive client/server applications, in *Proceedings of the IEEE Third Workshop on Multimedia Signal Processing (MMSP)*, IEEE Operations Center, Piscataway, N.J., 1999, pp. 389–394,.
15. D. Santa Cruz, T. Ebrahimi, J. Askelöf, M. Larsson, and C. Cristopoulos, JPEG 2000 still image coding versus other standards, *Proc. SPIE* **4115,** 446–454, SPIE, Bellingham, Washington (2000).
16. 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)*, vol. 2, Tampere University of Technology, Tampere, 2000, pp. 673–676.
17. D. Santa Cruz and T. Ebrahimi, An Analytical Study of the JPEG2000 Functionalities, in *Proc. of the IEEE International Conference on Image Processing (ICIP)*, vol. 2, IEEE Operations Center, Piscataway, N.J., 2000, pp. 49–52.
18. R.C. Gonzalez and R.E. Woods, *Digital Image Processing*, Addison-Wesley, Reading, MA, 1992.
19. K. Sayood, *Introduction to Data Compression*, Academic Press, 2nd edition, San Diego, CA, 2000.
20. S. Mallat, *A wavelet tour of signal processing*, Academic Press, New York, 1998.
21. R. Grosbois, D. Santa Cruz, J. Askelöf, 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, http://jj2000.epfl.ch.

Color Plate 18. Tiling artifacts: frog image together with its JPEG version at a compression rate of 1:90. (*Steingrímsson and Simon,* pp. 572–585)

Color Plate 19. Color shifts: sand image together with its JPEG version at a compression rate of 1:100. (*Steingrímsson and Simon,* pp. 572–585)

Color Plate 20. Original Lena image (top left) and its multiscale decompositions after one (top right), two (bottom left) and three (bottom right) decomposition steps using the biorthogonal 9/7 wavelet. (*Steingrímsson and Simon,* pp. 572–585)

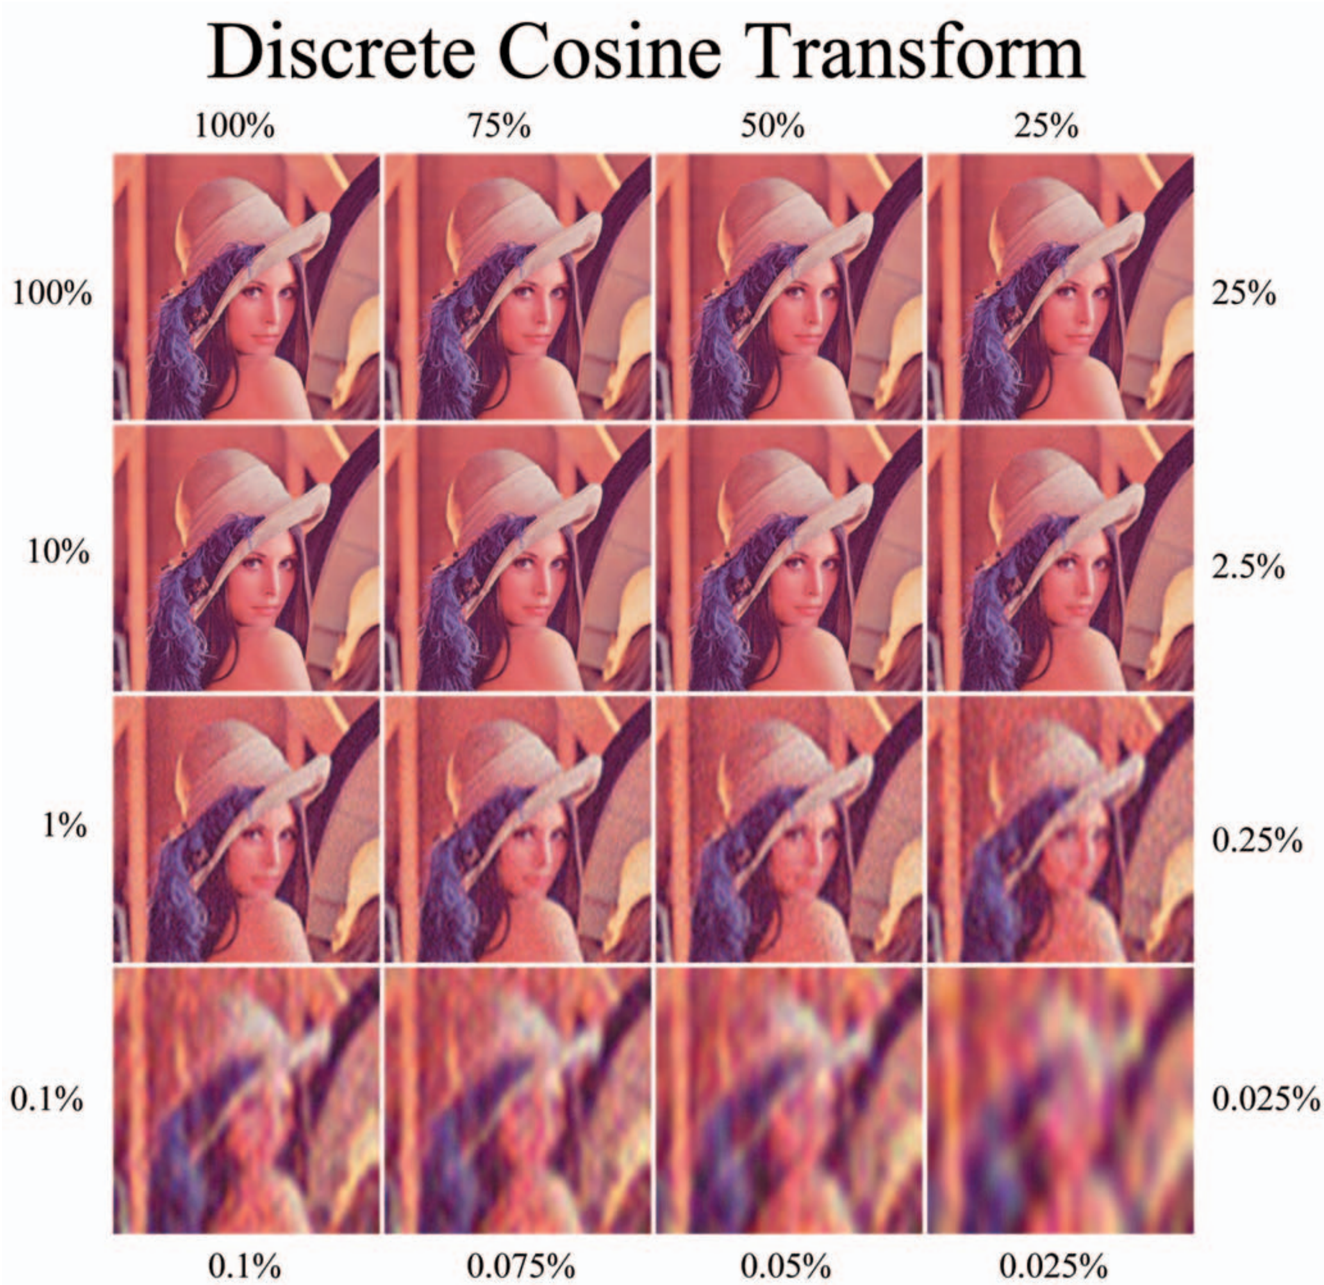

Color Plate 21. Full DCT transform and subsequent reconstruction using only 100%, 75%, 50%, 25%, 10%, 7.5%, 5%, 2.5%, 1%, 0.75%, 0.5%, 0.25%, 0.1%, 0.075%, 0.05% and 0.025% of the largest DCT coefficients, setting all remaining coefficients to zero. (*Steingrímsson and Simon,* pp. 572–585)

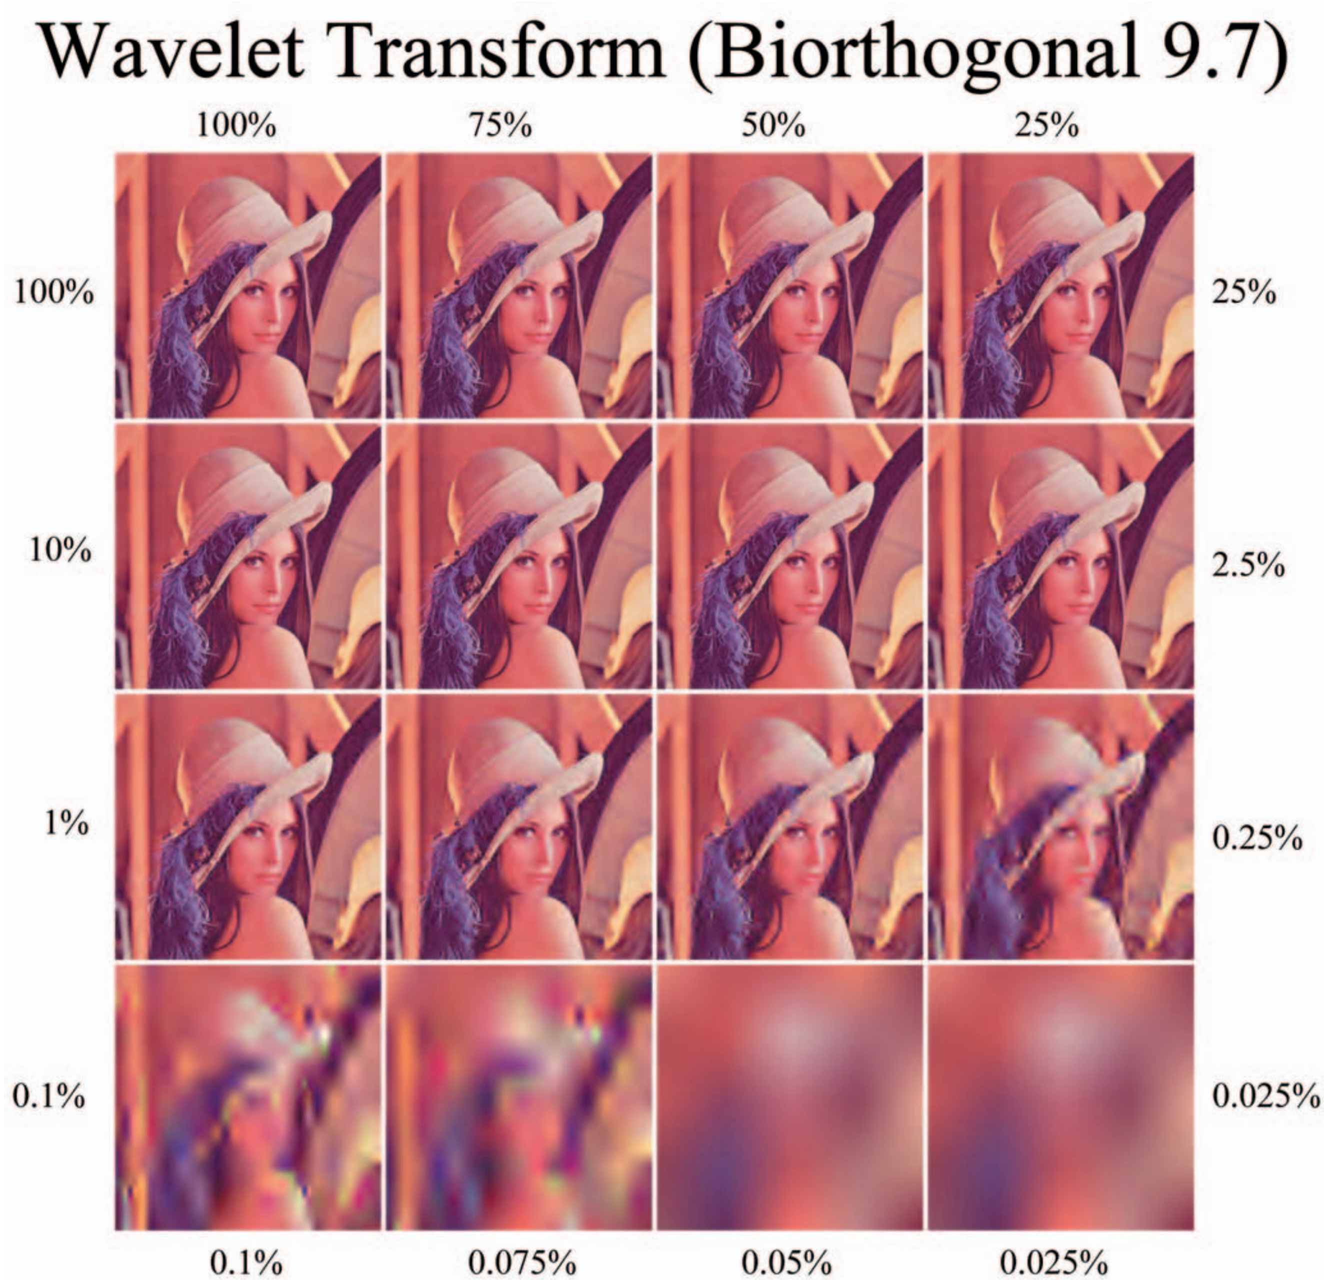

Color Plate 22. Full DWT transform with the biorthogonal 9/7 filter, and subsequent reconstruction using only 100%, 75%, 50%, 25%, 10%, 7.5%, 5%, 2.5%, 1%, 0.75%, 0.5%, 0.25%, 0.1%, 0.075%, 0.05% and 0.025% of the largest DWT coefficients, setting all remaining coefficients to zero. (*Steingrímsson and Simon,* pp. 572–585)

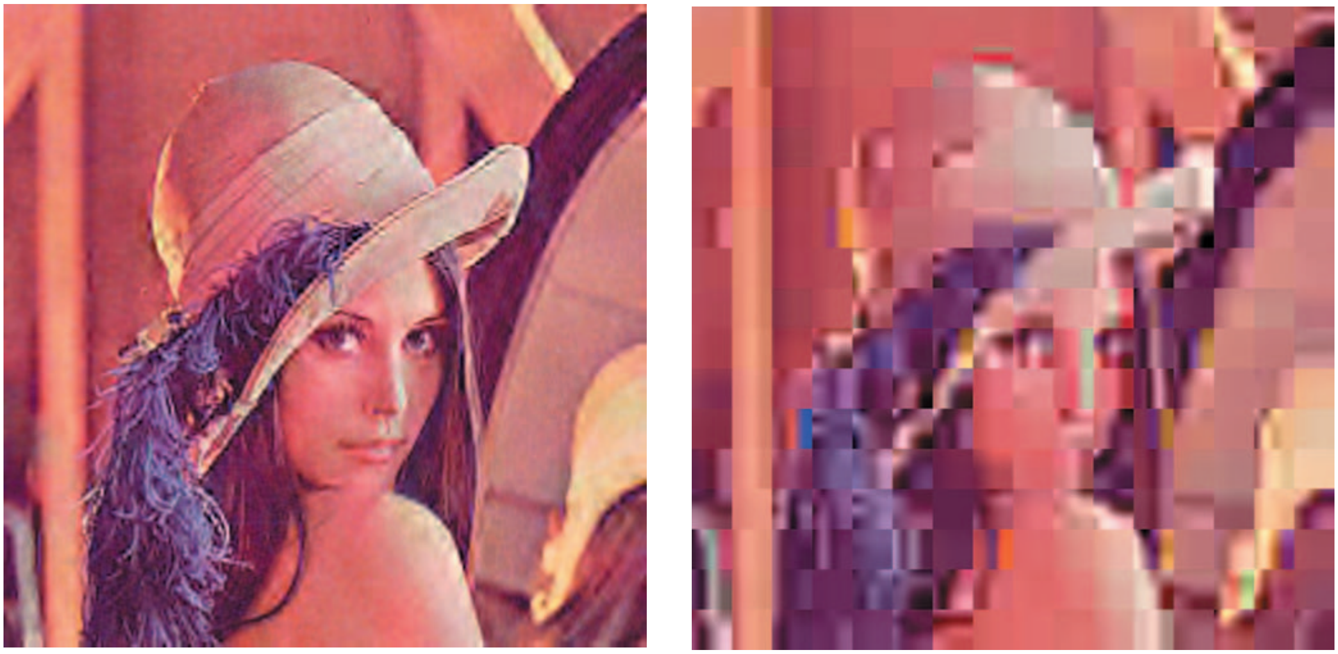

Color Plate 23. Tiling artifacts in JPEG 2000: compression of untiled (left) and tiled (32 × 32 pixel blocks, right) version of Lena, both at a compression rate of 1:48. (*Steingrímsson and Simon,* pp. 572–585)

Color Plate 24. Test images "Girl" (ISO 12640, top, left), "Cafe", (ISO 12640, top, right) and "Champagne", (ISO 12640, bottom). (*Steingrímsson and Simon,* pp. 572–585)

Color Plate 25. Large image collective forming the basis for the evaluations in section 'JPEG/JPEG 2000 quality matching'. The image resolution typically amounts to 1024 × 1280, approximately. (*Steingrímsson and Simon,* pp. 572–585)

Color Plate 26. Test layout for identifying the compression limit of visually perceivable distortions. This test example contains a JPEG compression affected by heavy distortions (center image), surrounded by two original versions (left and right). (*Steingrímsson and Simon,* pp. 572–585)

Color Plate 27. Test layout for quality comparisons: a JPEG (left) and a JPEG 2000 image (right) at random compression rates are confronted with each other. The examiner is now invited to select the image with fewer distortions with regard to the original. (*Steingrímsson and Simon,* pp. 572–585)

Color Plate 28. Images compressed at the rates emphasized with the bold cross and the bold circle of Fig. 6. The two images at the top display the JPEG 2000/JPEG image pair whose compression rates are illustrated by the bold cross and where the JPEG image on the right outperforms the JPEG 2000 image on the left with regard to reproduction authenticity. The image pair at the bottom exemplifies the inverse case which is indicated by the bold circle in Fig. 6. It corresponds to the case where the JPEG 2000 image on the left outperforms the JPEG image on the right. (*Steingrímsson and Simon,* pp. 572–585)

Color Plate 29. JPEG 2000 images produce typical “wavelet clouds” at higher compression rates (1:512). On the left is the original image, on the right the JPEG 2000 version.(*Steingrímsson and Simon,* pp. 572–585)

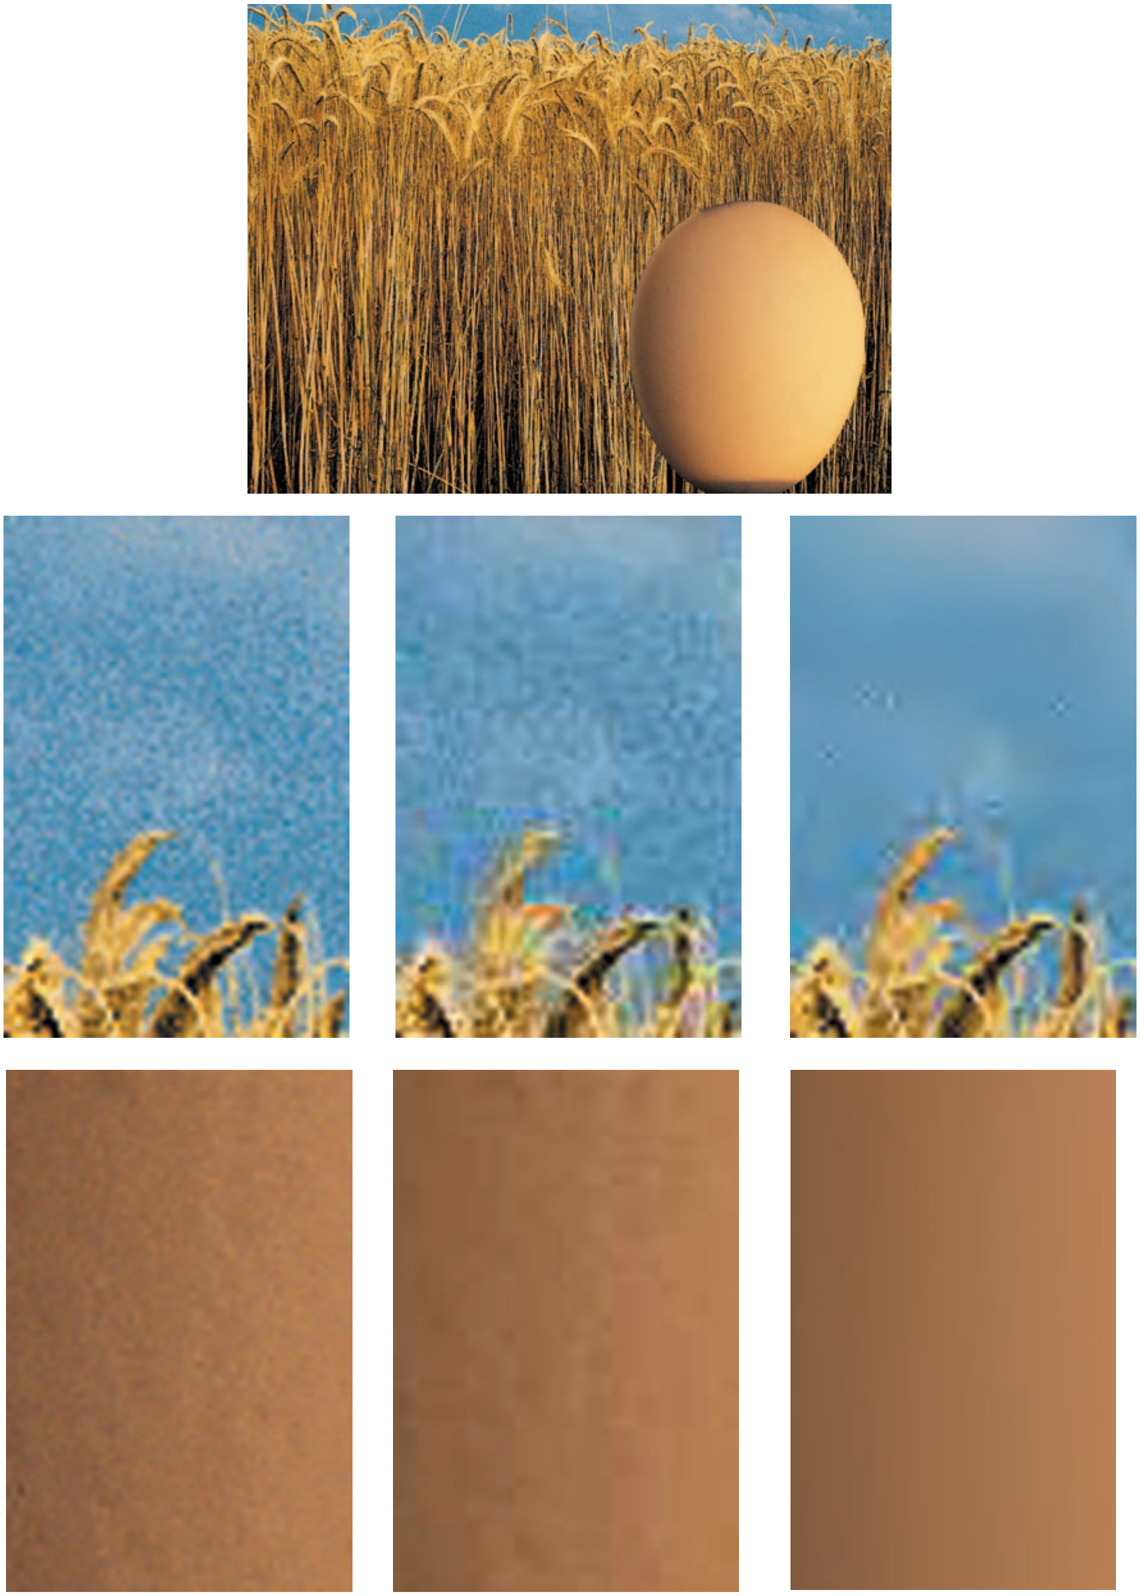

Color Plate 30. JPEG 2000 texture drop-off: at a compression rate of 12, the JPEG version of the original (top) image is still quite faithfully reproduced. The JPEG 2000 version at the same rate maps the high spatial frequencies of the corn field immaculately; the subtle granular structures on the egg surface and in the sky are, however, dramatically smoothed. Since distortions due to lossy compression are slight, they are, however, bothersome at high resolution; some magnified cutouts of the sky (middle image row) and of the egg surface (bottom image row) are displayed rather than the full images. The images on the left show the original sections. The JPEG version still preserves the smooth structures well (center images). In the JPEG 2000 version, however, the structures are almost completely missing (images on the right). (*Steingrímsson and Simon,* pp. 572–585)