JPEG2000 vs. JPEG from an image retrieval point of view

Gerald Schaefer
School of Computing and Technology
The Nottingham Trent University
Nottingham, United Kingdom
gerald.schaefer@ntu.ac.uk

Abstract

It is well known that JPEG2000 has a variety of advantages over its predecessor JPEG. JPEG2000 not only allows images to be coded with clearly better visual image quality, it also addresses a series of other issues. Unification of lossless and lossy compression modes, robustness to bit-errors to allow image transmission over noisy channels and provision of regions of interest (ROIs) are only some of those that have been incorporated into the new standard. In this paper we look at the JPEG2000 vs. JPEG debate from an image retrieval standpoint. While it seems evident that image compression will have a negative effect on the performance of retrieval algorithms our aim is to provide quantitative results of how severe this performance drop would be for JPEG2000 compression in comparison to standard JPEG encoding. Our results show that while high compression causes problems for retrieval of JPEG images, the retrieval performance of JPEG2000 images is almost independent of compression ratio.

1 Introduction

Every day thousands of new images are created: pictures taken with digital cameras, images scanned in using colour scanners, or stills captured from video sequences. Despite the ever-increasing hardware and communication technology capabilities we are still confronted with both limited disk space and limited bandwidth resources. Therefore images are normally stored in compressed form in order to utilise the available resources as effectively as possible. JPEG [13], a DCT based image compression algorithm is the current ISO standard for the encoding of still images. Despite its potential to provide reasonable image quality at medium compression rates, the demand for other, "better" image coding algorithms soon emerged.

Much effort and time has been spent on what finally became the new wavelet-based JPEG2000 image compression standard [10]. JPEG2000 not only allows images to be coded with clearly better visual image quality compared to its predecessor, it also addresses a series of other issues. Unification of lossless and lossy compression modes, robustness to bit-errors to allow image transmission over noisy channels and provision of regions of interest (ROIs) are only some of those that have been incorporated into the new standard. It is therefore not surprising that all these aspects have been rigorously tested and compared to the standard JPEG algorithm some evidence of which can be found in [5].

In this paper we look at the JPEG2000 vs. JPEG debate from an image retrieval standpoint. While it seems evident that image compression will influence the performance of retrieval algorithms this has only recently been investigated quantitatively [7]. It was confirmed that the application of image compression worsens retrieval effectiveness by a small but not negligible amount. As the experiments carried out in [7] were based on JPEG coded images, in here we set out to compare the performance drop that results from the use of JPEG coding with that resulting from the application of JPEG2000. In other words, we compare the image retrieval performance on a JPEG coded image database with that of JPEG2000 images.

To do so, we make use of the recently released UCID (Uncompressed Colour Image Database) dataset, an image database where all images were preserved in their uncompressed state [9]. In our experiments we use six widely used retrieval algorithms: colour histograms compared using L_1 [12] and L_2 [2] norms, colour moments [11], colour coherence vectors [4], auto-correlograms [3] and spatial-chromatic histograms [1]. Each of the UCID model and query images is compressed to a series of different compression rates by both JPEG and JPEG2000. Image retrieval is then performed for each possible combination of model/query compression rates. Results show that the performance drop arising from JPEG2000 compression is indeed significantly smaller than that stemming from JPEG based retrieval.

The rest of the paper is organised as follows: Section 2 briefly describes the retrieval methods that we used for the evaluation. Section 3 presents experimental results for the UCID dataset under JPEG and JPEG2000 compression while Section 4 concludes the paper.

2 Retrieval methods

The methods we evaluated in our experiments were colour histograms compared using L1 [12] and L2 [2] norms, colour moments [11], colour coherence vectors [4], auto-correlograms [3] and spatial-chromatic histograms [1]. In the following we provide a brief description of each of them.

2.1 Colour histograms - histogram intersection

Colour histograms - simply obtained by (uniformly) quantising the colour space counting the number of pixels that fall in each bin - were first used for image retrieval by Swain and Ballard [12]. As distance measure they introduced (the complement of) histogram intersection defined as

$$d_{\mathrm{HIS}}(I_1, I_2) = 1 - \sum_{k=1}^{N} \min(H_1(k), H_2(k)) \quad (1)$$

where H_1 and H_2 are the colour histograms of images I_1 and I_2, and N is the number of bins used for representing the histogram. It can be shown [12] that histogram intersection is equivalent to the L_1 norm and hence a metric. We used $8 \times 8 \times 8$ RGB histograms in our experiments.

2.2 Colour histograms - QBIC

An alternative to the L_1 norm is to use the Euclidean distance (L_2 norm) between two histograms. This approach was taken in the QBIC system [2] where they also addressed the problem of possible false negatives due to slight colour shifts by taking into account the similarity between separate histogram bins. This can be expressed in a quadratic form distance measure as

$$d_{\mathrm{QBIC}}(I_1, I_2) = (H_1 - H_2)A(H_1 - H_2)^T \quad (2)$$

where H_1 and H_2 are again the two colour histograms (in the form of a vector) and A is an $N \times N$ matrix containing the inter-bin distances. We used the Munsell colour space divided into 256 bins (16 for hue, 4 for chroma and value respectively) to generate these histograms.

2.3 Colour moments

Stricker and Orengo [11] used colour moments as a compact colour descriptor for CBIR. The n^{th} central (normalised) moment of a colour distribution is defined as

$$M^n(I) = \sqrt[n]{\frac{1}{N}(M^1(I) - c(x, y))^n} \quad (3)$$

with

$$M^1(I) = \frac{1}{N}\sum c(x, y) \quad (4)$$

where N is the number of pixels in an image and $c(x, y)$ describes the colour of the pixel at location (x, y). For our experiments we used the first three moments in the HSV colour space. The distance between two images is defined as the sum of absolute distances between their moments (L_1 norm)

$$d_{\mathrm{MNT}}(I_1, I_2) = \sum_{i=1}^{n} |M^i(I_1) - M^i(I_2)| \quad (5)$$

2.4 Colour coherence vectors

Pass and Zabih [4] introduced colour coherence vectors as a method of introducing spatial information into the retrieval process. Colour coherence vectors consist of two histograms: one histogram of coherent and one of non-coherent pixels. Pixels are considered to be coherent if they are part of a continuous uniformly coloured area and the size of this area exceeds some threshold τ where τ is usually defined as 1% of the overall area of an image. The L_1 norm is used as the distance metric between two colour coherence vectors

$$d_{\mathrm{CCV}}(I_1, I_2) = \sum_{k=1}^{N} [|H_1^c(k) - H_2^c(k)| + |H_1^s(k) - H_2^s(k)|] \quad (6)$$

where H_i^c and H_i^s and are the histograms of coherent and non-coherent (scattered) pixels respectively. In our implementation we first blurred the image using a 3 x 3 averaging filter and used 8 x 8 x 8 RGB bins for representing the histograms.

2.5 Colour correlograms

Another approach to incorporate information on the spatial correlation between the colours present in an image was proposed by Huang et al. [3]. They introduced the notation of colour correlograms (CCRs) defined as

$$\gamma_{c_i,c_j}^{(k)}(I) = \mathrm{PR}_{p_1 \in I_{c_i}, p_2 \in I}[p_2 \in I_{c_j}, |p_1 - p_2| = k] \quad (7)$$

with

$$|p_1 - p_2| = \max |x_1 - x_2|, |y_1 - y_2| \quad (8)$$

where c_i and c_j denote two colours and (x_k, y_k) denote pixel locations. In other words, given any colour c_i in the image, γ gives the probability that a pixel at distance k away is of colour c_j.

As full colour correlograms are expensive both in terms of computation and storage requirements, usually a simpler form called auto-correlogram (ACR) defined as

$$\alpha_c^{(k)}(I) = \gamma_{c,c}^{(k)}(I) \tag{9}$$

is being used, i.e. only the spatial correlation of each colour to itself is recorded. Two CCRs are compared using

$$d_{\text{CCR}}(I_1, I_2) = \frac{\sum_{i,j \in [m], k \in [d]} |\gamma_{c_i,c_j}^{(k)}(I_1) - \gamma_{c_i,c_j}^{(k)}(I_2)|}{\sum_{i,j \in [m], k \in [d]} (1 + \gamma_{c_i,c_j}^{(k)}(I_1) + \gamma_{c_i,c_j}^{(k)}(I_2))} \tag{10}$$

We used ACRs with $8 \times 8 \times 8$ RGB colours, for k we chose $\{1, 3, 5, 7\}$.

2.6 Spatial-chromatic histograms

Cinque et al. [1] introduced spatial-chromatic histograms (SCHs) as an alternative method for representing both colour and spatial information. SCHs consist of a colour histogram

$$h(k) = \frac{|A_k|}{n * m} \tag{11}$$

where A_k is a set having the same colour k, and n and m are the dimensions of the image; and location information on each colour characterised through its baricentre

$$b(k) = \left(\frac{1}{n} \frac{1}{|A_k|} \sum_{(x,y) \in A_k} x, \frac{1}{m} \frac{1}{|A_k|} \sum_{(x,y) \in A_k} y \right) \tag{12}$$

and the standard deviation of distances of a given colour from its baricentre

$$\sigma(k) = \sqrt{\frac{1}{|A_k|} \sum_{p \in A_k} d(p, b(k))^2} \tag{13}$$

The SCH is then given as

$$H_{\text{SCH}}(k) = [h(k), b(k), \sigma(k)] \tag{14}$$

and similarity between two SCHs calculated as

$$d_{\text{SCH}}(I_1, I_2) = 2 - \sum_{k=1}^{N} \min(h_{I_1}(k), h_{I_2}(k)) \cdot \\ \cdot \left(\frac{\sqrt{2} - d(b_{I_1}(k), d(b_{I_2}(k))}{\sqrt{2}} + \frac{\min(\sigma_{I_1}(k), \sigma_{I_2}(k))}{\max(\sigma_{I_1}(k), \sigma_{I_2}(k))} \right) \tag{15}$$

In our implementation we divided the Munsell colour space uniformly into 512 areas whose centres were used as the colours to describe spatial-chromatic histograms.

↓ Q/M →	-	100	80	50	30	20	10	5	0
-	95.27	95.20 95.30	95.25 95.31	95.15 95.28	95.16 95.26	94.98 95.26	94.21 95.22	91.24 95.24	85.53 95.26
100	95.16 95.30	95.15 95.29	95.17 95.30	95.16 95.26	95.25 95.26	95.08 95.26	94.52 95.24	91.91 95.30	85.45 95.23
80	95.14 95.16	95.13 95.19	95.15 95.28	95.18 95.29	95.28 95.27	95.08 95.21	94.52 95.22	92.01 95.31	85.52 95.23
50	94.99 95.12	95.05 95.09	95.07 95.26	95.09 95.26	95.17 95.24	94.98 95.19	94.46 95.19	92.14 95.26	85.72 95.28
30	95.00 95.09	95.08 95.05	95.12 95.16	95.17 95.22	95.19 95.22	95.01 95.24	94.47 95.22	92.02 95.35	85.90 95.32
20	94.84 95.02	95.03 94.94	95.06 95.14	95.09 95.21	95.17 95.18	94.98 95.23	94.50 95.22	92.31 95.29	85.72 95.28
10	94.74 94.81	95.02 94.78	95.12 95.05	95.14 95.05	95.22 95.08	94.89 95.18	94.84 95.21	92.98 95.20	85.91 95.19
5	89.25 94.67	89.25 94.65	89.35 94.95	89.45 95.04	89.52 95.04	89.42 95.08	87.51 95.03	94.55 95.10	86.08 95.13
0	77.24 94.57	76.65 94.54	76.47 94.91	76.11 95.01	76.23 95.05	75.67 95.07	74.53 95.02	77.87 95.10	92.40 95.12

Table 1. Retrieval performance for colour histograms. Top rows and outer left column indicate q-factors used for model and query images respectively ('-' indicates uncompressed images). Each cell contains AMP results for JPEG (top) and JPEG2000 (bottom) compressed images.

3 Experimental Results

We compressed each of the UCID images[1] to a series of different compression rates by using different JPEG q-factors: 100, 80, 50, 30, 20, 10, 5, and 0. The JPEG2000 images were then generated by simply matching the compression ratios achieved for the JPEG images (making use of another advantage of the JPEG2000 scheme, that of being able to precisely define target bit rates). For each possible query/model combination of compression rates (i.e. query images compressed with q=100, model images with q=100; queries with q=100, models with q=80; ...; query images with q=80, model images with q=100; ...; queries with q=0, models with q=0) image retrieval was performed (using the six methods outlined above) once with the JPEG images and once using the JPEG200 database. Results, expressed in terms of average match percentile (AMP)[2] for colour histograms, colour moments, and spatial-chromatic histograms are given in Tables 1, 2, and 3 respectively. Due to space restrictions we had to omit those for QBIC histograms, colour coherence vectors, and colour auto-correlograms. However the reader be assured that the observations made below for the results in the listed tables also hold for these other methods; also complete results can be found in a technical report [6]. Looking at these results we can first confirm what has been pointed out in [7], namely that for low to medium JPEG compression there is only a very slight drop in re-

[1]For the experiments presented here we used Version 1 of the UCID dataset [8] which comprises a subset of the database reported in [9]. Full results for both datasets can be found in [6]. Both versions of the UCID dataset are available from `http://vision.doc.ntu.ac.uk`.

[2]Since there is a 1:n correspondence between model and query images, the modified AMP from [9] was used. An AMP of 100 means perfect retrieval, while an AMP of 50 indicates that the systems is equivalent to one that returns images in a random order.

trieval performance. On the other hand however, compression with low to very low q-factors will produce results that are notably below those that are achieved for uncompressed (or slightly compressed) images. Inspecting the retrieval results that were obtained for JPEG2000 compressed images the picture is quite different. Indeed, it can be seen that the retrieval performance is almost independent of the compression with average match percentiles dropping by less than 1% for most cases! This performance is in fact better than we were hoping for before conducting the experiments and clearly demonstrates the superiority of JPEG2000 compared to its predecessor.

↓ Q/M →	-	100	80	50	30	20	10	5	0
-	92.29	91.38 92.35	90.75 91.91	90.47 91.41	89.89 91.21	88.52 90.85	80.84 90.25	70.80 89.57	68.87 88.70
100	91.83 92.28	92.16 92.31	91.86 91.86	91.88 91.32	91.94 91.18	91.04 90.81	83.52 90.11	73.85 89.47	71.61 88.68
80	91.79 91.70	92.46 91.66	92.39 92.13	92.36 92.04	92.51 91.93	91.70 91.71	84.08 91.22	74.24 90.76	72.16 90.17
50	91.52 91.33	92.01 91.42	92.29 92.16	92.45 92.23	92.68 92.25	92.21 91.93	84.69 91.67	74.32 91.23	71.97 90.84
30	91.15 91.23	91.51 91.18	91.84 91.93	92.17 92.19	92.74 92.18	92.44 91.92	85.75 91.52	74.21 91.11	71.64 90.67
20	90.73 90.93	91.04 90.93	91.53 91.86	91.88 92.18	92.58 92.13	92.84 91.95	87.40 91.53	74.54 91.16	71.10 90.75
10	85.02 90.68	84.79 90.77	85.08 91.77	86.36 92.10	87.69 92.16	89.80 91.90	91.87 91.74	77.11 91.56	71.09 91.19
5	76.93 90.43	77.54 90.38	77.80 91.32	78.77 91.77	79.77 91.88	81.50 91.77	86.23 91.53	89.13 91.34	80.99 91.00
0	74.20 90.27	75.20 90.28	75.35 91.19	75.82 91.50	76.53 91.60	77.94 91.54	82.22 91.38	89.00 91.30	88.06 91.02

Table 2. Retrieval performance for colour moments. The layout is the same as in Table 1.

↓ Q/M →	-	100	80	50	30	20	10	5	0
-	96.46	96.67 96.47	96.68 96.68	96.59 96.64	96.57 96.37	96.70 96.34	95.17 96.15	91.78 96.15	84.97 96.09
100	96.52 96.42	96.87 96.43	96.83 96.63	96.90 96.57	96.81 96.33	97.09 96.28	95.30 96.09	91.74 96.02	83.92 96.01
80	96.48 96.31	96.88 96.36	96.84 96.80	96.96 96.83	96.83 96.77	97.06 96.76	95.17 96.69	91.78 96.68	83.84 96.66
50	96.43 96.21	96.89 96.29	96.75 96.79	97.00 96.89	96.77 96.83	97.05 96.89	95.31 96.84	91.84 96.85	83.81 96.91
30	96.37 96.10	96.67 96.18	96.77 96.72	96.78 96.82	96.90 96.88	96.88 96.89	95.52 96.85	92.31 96.93	85.00 96.95
20	96.42 96.05	96.67 96.09	96.81 96.71	96.75 96.79	96.83 96.88	97.07 96.91	95.67 96.88	92.33 96.89	85.88 96.93
10	94.17 95.84	94.33 95.91	94.06 96.53	94.66 96.71	94.09 96.81	94.71 96.83	95.69 96.80	93.35 96.83	83.18 96.87
5	85.66 95.66	85.18 95.73	85.46 96.48	85.26 96.63	85.27 96.72	85.26 96.77	89.77 96.73	95.26 96.71	82.89 96.80
0	77.01 95.65	76.35 95.64	76.73 96.25	76.82 96.48	76.26 96.60	75.71 96.65	77.00 96.61	78.46 96.65	92.87 96.73

Table 3. Retrieval performance for spatial-chromatic histograms. The layout is the same as in Table 1.

4 Conclusions

An extensive series of experiments were conducted to explore how the recent JPEG2000 image compression standard performs in comparison to its predecessor JPEG. While several evaluations have been reported in the literature covering various aspects of compression algorithms, our main interest was to see how JPEG2000 does compared to JPEG when it comes to image retrieval of compressed images. In summary, the experiments carried out confirm what we set out to demonstrate, namely that JPEG2000 is not only a better compression algorithm in terms of image quality and robustness, but that this also translates to a superior method when it comes to its application for image retrieval.

Acknowledgements

The author would like to acknowledge the support of the Nuffield Foundation under grant number NAL/00703/G.

References

[1] L. Cinque, S. Levialdi, and A. Pellicano. Color-based image retrieval using spatial-chromatic histograms. In *IEEE Int. Conf. Multimedia Computing and Systems*, pages 969–973, 1999.

[2] C. Faloutsos, W. Equitz, M. Flickner, W. Niblack, D. Petkovic, and R. Barber. Efficient and effective querying by image content. *Journal of Intelligent Information Retrieval*, 3(3/4):231–262, 1994.

[3] J. Huang, S. Kumar, M. Mitra, W.-J. Zhu, and R. Zabih. Image indexing using color correlograms. In *IEEE Int. Conference Computer Vision and Pattern Recognition*, pages 762–768, 1997.

[4] G. Pass and R. Zabih. Histogram refinement for content-based image retrieval. In *3rd IEEE Workshop on Applications of Computer Vision*, pages 96–102, 1996.

[5] D. Santa-Cruz, R. Grosbois, and T. Ebrahimi. JPEG 2000 performance evaluation and assessment. *Signal Processing: Image Communication*, 17(1):113–130, 2002.

[6] G. Schaefer. JPEG2000 vs. JPEG from an image retrieval point of view. Technical report, School of Computing and Technology, The Nottingham Trent University, Nottingham, United Kingdom, 2004.

[7] G. Schaefer and M. Stich. On the influence of image compression on the performance of content based image retrieval. In *6th Int. Conference on VISual Information Systems*, pages 426–431, 2003.

[8] G. Schaefer and M. Stich. UCID - An Uncompressed Colour Image Database. Technical report, School of Computing and Technology, The Nottingham Trent University, Nottingham, United Kingdom, 2003.

[9] G. Schaefer and M. Stich. UCID - An Uncompressed Colour Image Database. In *Storage and Retrieval Methods and Applications for Multimedia 2004*, volume 5307 of *Proceedings of SPIE*, pages 472–480, 2004.

[10] A. Skodras, C. Christopoulos, and T. Ebrahimi. JPEG2000: the upcoming still image compression standard. In *Proc. 11th Portuguese Conference on Pattern Recognition (RECPA00D)*, pages 359–366, 2000.

[11] M. Stricker and M. Orengo. Similarity of color images. In *Conf. on Storage and Retrieval for Image and Video Databases III*, volume 2420 of *Proceedings of SPIE*, pages 381–392, 1995.

[12] M. Swain and D. Ballard. Color indexing. *Int. Journal Computer Vision*, 7(11):11–32, 1991.

[13] G. Wallace. The JPEG still picture compression standard. *Communications of the ACM*, 34(4):30–44, 1991.