

ON THE LIMITS OF PERCEPTUALLY OPTIMIZED JPEG

Francesca De Simone, Lutz Goldmann, Denis Filimonov, and Touradj Ebrahimi

Multimedia Signal Processing Group (MMSPG)
Ecole Polytechnique Fédérale de Lausanne (EPFL)
CH-1015 Lausanne, Switzerland

ABSTRACT

The study presented in this paper aims at quantifying the empirical limits of JPEG optimization, when the compressed stream is standard compliant and only the quantization tables are optimized. Image-dependent quantization tables, which minimize the bitrate of the compressed image while maintaining transparent visual quality, are identified by means of a psychovisual experiment. The results demonstrate that significant room for improving JPEG compression efficiency is available.

1. INTRODUCTION

Today, several alternatives for compression of digital pictures exist to choose from. Beside the well-known and largely deployed JPEG compression, other approaches, standardized by international committees, include JPEG 2000, H.264/AVC Intra, and the recently adopted JPEG XR. With respect to JPEG compression, the other solutions have been claimed to produce a significant compression gain. This is mainly due to the usage of more efficient transforms, optimized quantization, and advanced entropy coding strategies.

Focusing on the quantization strategy, many algorithms for optimizing this step of the coding chain have been proposed in literature. Some of them are based on the minimization of a distortion measure, like the block-based optimization by Watson [1] and the image-adaptive strategy by Fung *et al.* [2] for JPEG compression, and the recent works by Richter [3] and Schonberg *et al.* [4] for JPEG XR compression. Others rely on models of the human visual system, like the work by Battiato *et al.* [5] for JPEG compression and the work by Nadenau [6] for JPEG 2000 compression. On the other hand, at the best of the authors' knowledge, no studies are available which try to quantify the maximum room for compression gain which is attainable by subjectively optimizing the quantization tables.

This work was supported by the Swiss National Foundation for Scientific Research in the framework of NCCR IM2 and by the European Network of Excellence Petamedia.

The study presented in this paper aims at quantifying the empirical limits of JPEG optimization, when the compressed stream is standard compliant and only the quantization tables are optimized. Particularly, image-dependent quantization tables, which minimize the bitrate of the compressed image while maintaining transparent visual quality, are identified by means of a psychovisual experiment. The optimization procedure has been applied separately to each image, considering a set of six high resolution images having different color and texture features.

Using the identified image-dependent tables, gains of 15% to 22% in terms of compression efficiency can be reached with respect to the image-independent quantization tables specified in Annex K of the standard [7]. As discussed in Section 3, the tables resulting from our experiment are not meant to be the "optimal" tables, i.e. the tables, among all the possible combinations of 8x8 values of quantization step size, which allow achieving the minimum possible bitrate for transparent quality of the coded image. On the other hand, they clearly show that compression gain of at least 15% to 20% can be achieved by using optimized quantization tables, thus, significant room for improvement of JPEG compression efficiency is still available.

The rest of the paper is organized as follows: an overview of the quantization strategy used in the JPEG standard is provided in Section 2; the method for the experimental optimization of JPEG quantization tables is described in Section 3; finally results and conclusions are discussed in Sections 4 and 5, respectively.

2. QUANTIZATION TABLES IN JPEG

Figure 1 shows the overall JPEG image compression architecture. In this section we will focus on the analysis of the quantization strategy, thus, please refer to [7] for further details regarding the other steps of the compression scheme.

The image is first divided into blocks of 8x8 pixels, on which the forward discrete cosine transform (FDCT) is applied, followed by quantization and entropy coding. After the FDCT, each of the 64 DCT coefficients is uniformly quantized in conjunction with a 64-element quantization ta-

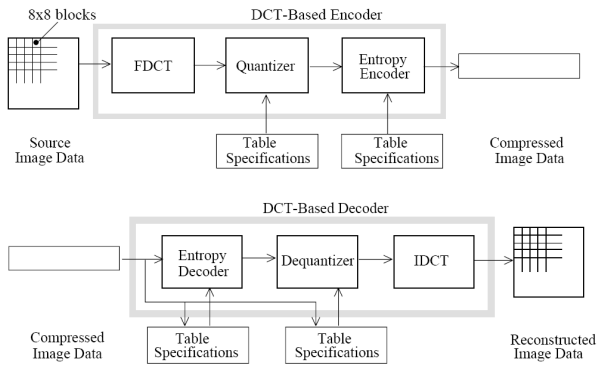


Fig. 1: DCT-based encoder and decoder processing steps in JPEG compression [7].

ble, according to the following rule:

$$DCT'(u, v) = \text{int} (DCT(u, v) / \Delta(u, v)) \quad (1)$$

where $u = \{1, \dots, 8\}$ and $v = \{1, \dots, 8\}$, $DCT(u, v)$ is the DCT coefficient in the u -th row and v -th column of the 8×8 DCT block, $\Delta(u, v)$ is the corresponding step size of the quantizer specified in the 8×8 quantization table, and $DCT'(u, v)$ is the quantized DCT coefficient.

The quantization table must be specified as an input to the encoder. At the decoder, it is extracted from the data stream. After entropy decoding, the dequantized values of DCT coefficients are obtained by multiplying each decoded value to the corresponding value in the quantization table.

For baseline JPEG, each element in the quantization table can be any integer value in the range $[1, 255]$. The compression rate can be adjusted by varying the quantization step sizes, allowing a tradeoff between storage size or bitrate and image quality.

Typically, two different tables are used for the luminance component and for the chrominance components of the image. The Annex K of the standard provides as an example the quantization tables depicted in Table 1. These tables have been the results of psychovisual experiments performed to investigate upon the visibility of 8×8 DCT basis functions [8].

To allow a different quality of the output image to be selected by the user, a common practice is to scale the quantization tables by setting a parameter which ranges from 1 to 100, usually called "quality factor" ($Qfactor$). In the widely used IJG free library for JPEG image compression [9], the final quantization table is obtained by scaling the input 8×8 table according to the input quality factor:

$$\Delta(u, v) = \text{int} \left(\frac{Qtable(u, v) \cdot Sfactor + 50}{100} \right) \quad (2)$$

Luminance								Chrominance							
16	11	10	16	24	40	51	61	17	18	24	47	99	99	99	99
12	12	14	19	26	58	60	55	18	21	26	66	99	99	99	99
14	13	16	24	40	57	69	56	24	26	56	99	99	99	99	99
14	17	22	29	51	87	80	62	47	66	99	99	99	99	99	99
18	22	37	56	68	109	103	77	99	99	99	99	99	99	99	99
24	35	55	64	81	104	113	92	99	99	99	99	99	99	99	99
49	64	78	87	103	121	120	101	99	99	99	99	99	99	99	99
72	92	95	98	112	100	103	99	99	99	99	99	99	99	99	99

Table 1: Quantization tables for the luminance and chrominance components, defined in Annex K of the JPEG standard [7].

and

$$\Delta(u, v) = \begin{cases} 1 & \text{if } \Delta(u, v) \leq 0 \\ 255 & \text{if } \Delta(u, v) > 255 \end{cases} \quad (3)$$

with $Qtable(u, v)$ denoting the element in the u -th row and v -th column in the input table, and $Sfactor$ a scaling factor which depends upon the value of $Qfactor$:

$$Sfactor = \begin{cases} \text{int}(5000/Qfactor) & \text{if } Qfactor < 50 \\ 200 - 2 \cdot Qfactor & \text{otherwise} \end{cases} \quad (4)$$

When $Qfactor$ is equal to 100, the coding is lossless, as the final quantization step sizes are all equal to 1. It can be noticed that, due to the round to integer operation in (1) and (3) and to the saturation operation in (2), for a fixed value of $Qfactor$, not all the values in the range $[1, 255]$ for $Qtable(u, v)$ are significant. In other words, for a fixed $Qfactor$, many different values of $Qtable(u, v)$ lead to the same $\Delta(u, v)$.

3. EXPERIMENTAL OPTIMIZATION

As proven in many studies [1] [2] [5], the quantization tables in Table 1 are "sub-optimal", since they are independent from the features of the particular image under analysis. Theoretically speaking, the optimal tables, i.e. which allows to achieve the minimum possible bitrate for a certain quality of the coded image, could be identified, by performing an exhaustive search over the set of all possible tables for each image. As the values of the step sizes are in the range $[1, 255]$, there are 255^{64} possible quantization tables, for each component of the image. Obviously, this kind of search is not feasible.

Since our study focuses on the identification of optimal tables for transparent quality of the coded image, one approach to reduce the dimensionality of the problem could be based on the IJG quality scaling strategy. As described in Section 2, once identified the $Qfactor$ which assures transparent quality of the coded image produced by using the recommended tables, the range of significant values for the

quantization step sizes would be reduced. Thus, the number of different combinations of values in the tables would be reduced as well. This option has been investigated but the final ranges of possible values were still too large to make the approach of extensive search feasible. Therefore, the procedure described below has been applied. The IJG software has been used to produce all the test material, using baseline JPEG coding.

First, for each image separately, an experiment has been performed in order to identify the value of $Qfactor$ which assures transparent quality when using the recommended tables, according to the following steps:

- The entire set of compressed pictures using the recommended tables has been created by varying $Qfactor$ from 1 to 100.
- A subset of the compressed images has been selected, by excluding all the samples which presented strong quality degradations and ordering the remaining samples from the lowest to the highest bitrate.
- An expert observer selected the compressed image in the set of ordered samples, which minimizes the bitrate while presenting no perceptual difference in relation to the original image. To facilitate this task an interface has been developed, which shows the original image on one side of the screen and the compressed image on the other side, allowing the user to sequentially browse the compressed images in the set starting from the lowest bitrate.

After this first step, a second experiment has been carried out. In this experiment, an expert viewer was free to modify the elements of the quantization tables. Starting from the recommended tables, the user's goal was to identify new tables which allow for a transparent quality of the coded image while reducing its bitrate.

In order to identify optimized image-dependent quantization table, a Graphical User Interface (GUI) has been developed which allows the user to interactively modify the coefficients of the luminance and the chrominance quantization tables used in the JPEG encoding/decoding process. Considering an input test picture and fixing the $Qfactor$ identified with the previous experiment, the GUI displays the coded image produced by using the currently defined tables, side by side to the original uncompressed and to the compressed version obtained by using the image-independent tables specified in Annex K of the standard. The bitrates of the two compressed pictures are also displayed. This way, the user can immediately observe the effect of the modifications and change the combinations of values in the quantization tables. A screenshot of the GUI is shown in Figure 2.

To validate the results of the optimization procedure, the pictures selected by one expert viewer have been inspected by an additional group of five expert viewers. A small test was performed including the original uncompressed picture, the picture compressed using the recommended tables and transparent $Qfactor$, and the pictures compressed with the transparent $Qfactor$ but using the optimized tables. The pictures were shown in pairs, presenting all the possible 3^2 combinations, randomly mixed. None of the five viewers could detect any differences among the stimuli, i.e. none of the subjects was able to distinguish which one of the three images was the original.

A test set of six high resolution (1280x1600 pixels) color images having different color and texture features has been considered (Figure 3). The uncompressed images have 24 bit per pixel and are available for download at <http://mmspg.epfl.ch/iqa>. Due to the extremely high resolution of the data, a test set-up with two LCD monitors with native resolution of 2560x1600 pixels has been used. The GUI commands and the original picture, at its native resolution, have been displayed on the first monitor; the decoded picture produced using the recommended tables and the decoded picture produced using the current tables have been displayed, at their native resolution, on the second monitor. The monitors have been connected to the same server and calibrated using an EyeOne Display2 color calibration device according to the following profile: sRGB Gamut, D65 white point, 120 cd/m^2 brightness and minimum black level. The ambient lighting system consisted of neon lamps with 6500 K color temperature. The illumination level measured on the screen was 30 lux and the ambient black level was 0.5 cd/m^2 .



Fig. 3: Screenshots of the six test pictures. From top left, clockwise: $p01$, $p06$, $p10$, $bike$, $cafe$, $woman$.

4. RESULTS AND DISCUSSION

Table 2 shows the results of the experiments described in the previous section, for each test image: the value of $Qfactor$ which assures transparent quality of the image coded using the recommended tables, with the corresponding bit per




Fig. 2: Screenshot of the GUI used to perform the optimization of quantization tables.

pixel value - bpp (rec. tables); the bit per pixel value of the image coded using the new tables - bpp (ad hoc tables); the gain achieved using the new tables, in terms of bit per pixel saving. Table 3 also shows the Peak Signal to Noise Ratio and the Structural Similarity index values [10] computed for the luminance and the chrominance components of the compressed pictures, before and after the optimization of the quantization tables. The optimized quantization tables are reported in Appendix, together with the final tables obtained by applying equation (1).

Image	Qfactor	bpp (rec. tables)	bpp (ad hoc tables)	bpp gain(%)
p01	87	1.93	1.60	16.8%
p06	92	2.16	1.67	22.7%
p10	80	1.39	1.19	14.3%
bike	84	2.06	1.76	14.5%
cafe	95	5.85	4.97	14.2%
woman	81	1.70	1.35	20.4%

Table 2: Results of the experiments described in Section 3.

The results agree with the findings of other studies [2], showing that a gain of 15% to 20% in JPEG compression

Image	PSNR Y, Cb, Cr (dB)		SSIM Y, Cb, Cr	
	Rec. tables	Ad hoc tables	Rec. tables	Ad hoc tables
p01	42.8;48.5;41.9	38.0;48.3;41.9	0.99;0.99;0.96	0.97;0.99;0.96
p06	44.7;41.7;42.6	39.6;41.7;42.5	0.98;0.96;0.97	0.97;0.96;0.97
p10	42.0;46.4;47.0	40.5;46.4;47.0	0.98;0.99;0.99	0.98;0.99;0.99
bike	38.5;40.2;41.0	37.9;40.2;41.0	0.96;0.93;0.95	0.95;0.93;0.95
cafe	42.5;34.5;35.9	36.8;34.4;35.8	0.99;0.88;0.91	0.98;0.88;0.91
woman	37.9;44.4;43.4	33.8;44.4;43.3	0.96;0.97;0.96	0.93;0.97;0.96

Table 3: PSNR and SSIM values before and after the optimization of quantization tables.

performance can be reached by simply optimizing the quantization tables. It is interesting to notice that the images for which the maximum gain is reached (around 20%), are *p06* and *woman*, which both contain large flat areas. Thus, the high frequencies can be quite strongly quantized without affecting the visual quality. This is also confirmed by the final quantization tables, reported in Table 11 and 15. Additionally, for images *p06*, *p10* and *woman*, no modifications in the quantization tables of the chrominance components have been made. This can be explained by the fact that these images contains flat areas with uniform color, thus, if a coarse quantization of the chrominance components occurs, artifacts in the chrominance components are easily noticeable. Finally, considering the PSNR and SSIM values in Table 3, it is evident that, in this case, both metrics are not always correctly predicting the perceived subjective quality: five expert viewers who accurately inspected and compared the images coded with the recommended and the optimized tables did not detect any perceivable difference.

5. CONCLUSIONS

The study presented in this paper aims at quantifying the empirical limits of JPEG optimization, when the compressed stream is standard compliant and only the quantization tables are optimized. Image-dependent quantization tables, which minimize the bitrate of the compressed image while maintaining transparent visual quality, are identified by means of a psychovisual experiment. The results obtained over a set of six different high resolution images show that compression gains up to 22% can be reached by using the image-dependent optimized quantization tables, with respect to image independent tables recommended in the JPEG standard.

It is important to underline that the tables resulting from our experiment are not meant to be the "optimal" tables, i.e. the tables, among all the possible combinations of 64 values of quantization step size, which allow achieving the minimum possible bitrate for transparent quality of the coded image. On the other hand, they clearly show that there is still a lot of room for improvement of the compression efficiency of the well established JPEG images compression standard, by taking into account image-dependent perceptual optimization strategies.

REFERENCES

- [1] A. B. Watson, "DCT quantization matrices visually optimized for individual images," *Proc. SPIE Human Vis., Visual Proc., Digital Disp. IV*, pp. 1913–14, 1993.
- [2] H. T. Fung and K. J. Parker, "Design of image-adaptive quantization tables for JPEG," *Journal of Electronic Imaging*, vol. 4(2), pp. 144–150, 1995.
- [3] T. Richter, "Visual quality improvement techniques of HD Photo/JPEG XR," in *Proc. Int. Conf. Image Proc.*, 2008.
- [4] D. Schonberg, G. Sullivan, S. Sun, and Z. Zhou, "Perceptual encoding optimization for JPEG XR image coding using spatially adaptive quantization step size control," *Proc. SPIE*, vol. 7443, 2009.
- [5] S. Battiato, M. Mancuso, A. Bosco, and M. Guarnera, "Psychovisual and statistical optimization of quantization tables for DCT compression engines," *Proc. Int. Conf. Image Analysis and Proc.*, pp. 602–606, 2001.
- [6] M. Nadenau, *Integration of human color vision models into high quality image compression.*, Ph.D. thesis, EPFL, 2000.
- [7] ISO/IEC IS 10918-1 ITU-T Recommendation T.81, "Digital compression and coding of continuous still images," 1991.
- [8] H. Lohscheller, "A subjectively adapted image communication system," *IEEE Trans. Commun. COM 32*, pp. 1316–1322, 1984.
- [9] Independent JPEG Group library release 7, "http://www.ijg.org/," 2009.
- [10] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE Trans. Im. Proc.*, vol. 13(4), pp. 600–612, 2004.

ANNEX

Luminance								Chrominance							
18	15	15	16	24	255	255	255	30	18	24	47	99	99	99	99
15	15	14	19	26	255	255	255	8	21	26	66	99	99	99	99
15	13	16	24	255	255	255	255	24	26	56	99	99	99	99	99
14	17	22	255	255	255	255	255	47	66	99	99	99	99	99	99
255	255	255	255	255	255	255	255	99	99	99	99	99	99	99	99
255	255	255	255	255	255	255	255	99	99	99	99	99	99	99	99
255	255	255	255	255	255	255	255	99	99	99	99	99	99	99	99
255	255	255	255	255	255	255	255	99	99	99	99	99	99	99	99
255	255	255	255	255	255	255	255	99	99	99	99	99	99	99	99

Table 4: Optimized quantization table for image *p01*.

Luminance								Chrominance							
25	25	30	16	24	40	255	255	20	18	24	47	99	99	99	99
30	25	25	19	26	255	255	255	18	21	26	66	99	99	99	99
20	25	16	24	255	255	255	255	24	26	56	99	99	99	99	99
14	17	22	255	255	255	255	255	47	66	99	99	99	99	99	99
18	22	255	255	255	255	255	255	99	99	99	99	99	99	99	99
24	255	255	255	255	255	255	255	99	99	99	99	99	99	99	99
255	255	255	255	255	255	255	255	99	99	99	99	99	99	99	99
255	255	255	255	255	255	255	255	99	99	99	99	99	99	99	99
255	255	255	255	255	255	255	255	99	99	99	99	99	99	99	112

Table 5: Optimized quantization table for image *p06*.

Luminance								Chrominance							
15	20	20	20	25	40	255	255	17	18	24	47	99	99	99	99
20	20	20	25	25	40	255	255	18	21	26	66	99	99	99	99
20	20	25	25	50	70	255	255	24	26	56	99	99	99	99	99
20	25	25	50	60	255	255	255	47	66	99	99	99	99	99	99
25	25	50	60	255	255	255	255	99	99	99	99	99	99	99	99
40	40	70	255	255	255	255	255	99	99	99	99	99	99	99	99
255	255	255	255	255	255	255	255	99	99	99	99	99	99	99	99
255	255	255	255	255	255	255	255	99	99	99	99	99	99	99	99

Table 6: Optimized quantization table for image *p10*.

Luminance								Chrominance							
15	20	20	20	25	40	255	255	17	18	24	47	99	99	99	99
20	20	20	25	25	40	255	255	18	21	26	66	99	99	99	99
20	20	25	25	50	70	255	255	24	26	56	99	99	99	99	99
20	25	25	50	60	255	255	255	47	66	99	99	99	99	99	99
25	25	50	60	255	255	255	255	99	99	99	99	99	99	99	99
40	40	70	255	255	255	255	255	99	99	99	99	99	99	99	99
255	255	255	255	255	255	255	255	99	99	99	99	99	99	99	99
255	255	255	255	255	255	255	255	99	99	99	99	99	99	99	99

Table 7: Optimized quantization table for image *bike*.

Luminance								Chrominance							
20	20	10	16	24	40	51	255	25	18	24	47	99	99	99	99
20	20	20	19	26	58	255	255	18	21	26	66	99	99	99	99
14	20	16	24	40	255	255	255	24	26	56	99	99	99	99	99
14	17	22	29	51	255	255	255	47	66	99	99	99	99	99	99
18	22	37	56	255	255	255	255	99	99	99	99	99	99	99	99
24	35	55	255	255	255	255	255	99	99	99	99	99	99	99	99
49	255	255	255	255	255	255	255	99	99	99	99	99	99	99	99
255	255	255	255	255	255	255	255	99	99	99	99	99	99	99	99

Table 8: Optimized quantization table for image *cafe*.

Luminance								Chrominance							
16	16	16	20	25	40	51	61	17	18	24	47	99	99	99	99
16	16	16	25	60	58	255	255	18	21	26	66	99	99	99	99
16	16	25	30	30	255	255	255	24	26	56	99	99	99	99	99
20	25	30	30	255	255	255	255	47	66	99	99	99	99	99	99
25	60	30	255	255	255	255	255	99	99	99	99	99	99	99	99
40	35	255	255	255	255	255	92	99	99	99	99	99	99	99	99
49	255	255	255	255	255	255	255	99	99	99	99	99	99	99	99
72	255	255	255	255	100	255	255	99	99	99	99	99	99	99	99

Table 9: Optimized quantization table for image *woman*.

