

JPEG COMPRESSION HISTORY ESTIMATION FOR COLOR IMAGES

Ramesh Neelamani, Ricardo de Queiroz, Zhigang Fan, and Richard Baraniuk

neelsh@rice.edu, queiroz@ieee.org, zfan@crt.xerox.com, richb@rice.edu

ABSTRACT

We routinely encounter digital color images that were previously JPEG-compressed. We aim to retrieve the various settings—termed JPEG compression history (CH)—employed during previous JPEG operations. This information is often discarded en-route to the image’s current representation. The discrete cosine transform coefficient histograms of previously JPEG-compressed images exhibit near-periodic behavior due to quantization. We propose a statistical approach to exploit this structure and thereby estimate the image’s CH. Using simulations, we first demonstrate the accuracy of our estimation. Further, we show that JPEG recompression performed by exploiting the estimated CH strikes an excellent file-size versus distortion tradeoff.

1. INTRODUCTION

JPEG is a commonly used standard to compress digital color images [1]. It achieves compression by quantizing the discrete cosine transform (DCT) coefficients of the image’s three color planes; see Fig. 1 for an overview of JPEG. However, the various settings used during JPEG compression and decompression are not standardized [1]. The following JPEG settings can be chosen by the user such as an imaging device: (1) the color space used to independently compress the image’s three color planes; (2) the subsampling employed on each color plane during compression and the complementary interpolation used during decompression; and (3) the quantization table used to compress each color plane. We refer to the settings used during JPEG operations as the image’s *JPEG compression history* (CH).

An image’s CH is often not directly available from its current representation. For example, JPEG images are also routinely converted to lossless-compression formats such as bitmap (BMP) format (say, to create a background image for Windows or to feed a printing driver) or Tagged Image File Format (TIFF). In such cases, the JPEG compression settings are discarded after decompression. The CH, if available, can be used applications such as effective recompression of BMP and TIFF images; JPEG-compressing the image with previous JPEG settings yields significant file-size reduction without introducing any additional distortion. An image’s JPEG CH can also potentially be used as an authentication feature, for covert messaging, or to uncover the compression settings used inside digital cameras. Hence, the problem of CH estimation is useful.

The CH estimation problem is relatively unexplored. Fan and de Queiroz have proposed a statistical framework to perform CH estimation for grayscale images [2]; for a grayscale image, the CH comprises only the quantization table employed during previous

JPEG operations. For some special cases, Neelamani, de Queiroz, and Baraniuk performed CH estimation for color images by exploiting the lattice structure of quantized DCT coefficients via novel lattice-based algorithms [3]. However, the general problem of CH estimation for color images remains largely unsolved.

We propose a statistical framework to perform CH estimation. We realize that due to JPEG’s quantization operation, the DCT coefficient histograms of previously JPEG-compressed images exhibit near-periodic structure. We statistically characterize this near-periodicity for a single color plane. The resulting framework can be exploited to estimate a grayscale image’s CH, namely, its quantization table. We extend the statistical framework to color images and design a dictionary-based CH estimation algorithm that provides the *maximum a priori* (MAP) estimate of a color image’s CH

$$\{G^*, S^*, Q^*\} = \arg \max_{G, S, Q} P(\text{Image}, G, S, Q), \quad (1)$$

with G^* , S^* , Q^* denoting the estimated compression color space, the subsampling and associated interpolation, and the quantization tables.

Our proposed CH estimation algorithm demonstrates good performance in practice. Further, we verify that the estimated CH allows us to recompress an image with minimal distortion (large signal-to-noise-ratio (SNR)) and simultaneously achieve a small file-size.

2. JPEG OVERVIEW

In this section, we review JPEG compression and decompression [1]. Consider a color image that is currently represented in the F color space (see Fig. 1); $F1$, $F2$, and $F3$ denote the three color planes. We refer to the F space as the *observation color space*. Assume that the image was previously JPEG-compressed in the G color space—termed *compression color space*. JPEG compression performs the following operations independently on each color plane $G1$, $G2$, and $G3$ in the G space:

1. Optionally downsample each color plane (for example, retain alternate pixels to downsample by a factor of two); this process is termed *subsampling*.
2. Split each color plane into 8×8 blocks. Take the DCT of each block.
3. Quantize the DCT coefficients at each frequency to the closest integer multiple of the quantization step-size corresponding to that frequency. For example, if X denotes an arbitrary DCT coefficient and q denote the quantization step-size for the corresponding DCT frequency, then the quantized DCT coefficient \bar{X}_q is obtained by

$$\bar{X}_q := \text{round} \left(\frac{X}{q} \right) q. \quad (2)$$

R. Neelamani and R. Baraniuk are with the ECE Dept., Rice University, Houston, TX 77030, USA; R. de Queiroz is with the EE Dept., University of Brazil, CP 04591, Brasilia, DF, 70910-900, Brazil; and Z. Fan is with Xerox Corporation, 800 Phillips Road, Webster, NY 14580, USA.

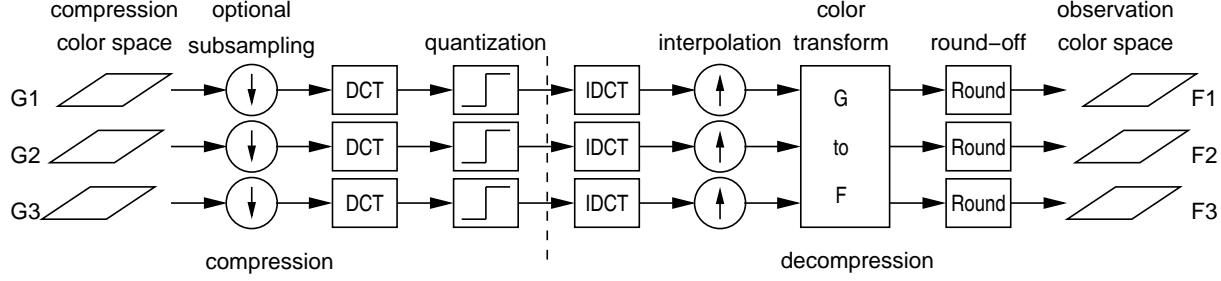


Fig. 1. Overview of JPEG compression and decompression.

JPEG decompression performs the following operations:

1. Take the inverse DCTs of the 8×8 blocks of quantized coefficients.
2. Interpolate the downsampled color planes by repetition followed by optional spatial smoothing with a low-pass filter. The popular IJG JPEG implementation [4] uses a $\frac{1}{4} \times [1 \ 2 \ 1]$ impulse response filter to smooth in the horizontal and vertical directions.
3. Transform the decompressed image to the desired color space F using the appropriate G to F transformation.
4. Round-off resulting pixel values to the nearest integer so that they lie in the 0–255 range.

3. CH ESTIMATION FOR GRAYSCALE IMAGES

For grayscale images, JPEG compression and decompression replicates the steps outlined in Section 2 for a single color plane but without subsampling and interpolation. Due to JPEG's quantization operations, the discrete cosine transform (DCT) coefficient histograms of previously JPEG-compressed grayscale images exhibit a near-periodic structure with the period determined by the quantization step-size. In this section, we derive a statistical framework, which characterizes the near-periodic structure, to estimate the quantization table.

An arbitrary DCT coefficient \tilde{X} of a previously JPEG-compressed grayscale image can be obtained by adding to the corresponding quantized coefficient \bar{X}_q (refer (2)) a round-off error term Γ

$$\tilde{X} = \bar{X}_q + \Gamma. \quad (3)$$

As described in [2], we can model Γ using a truncated Gaussian distribution

$$P(\Gamma = t) = \frac{1}{\Upsilon} \exp\left(-\frac{t^2}{2\sigma^2}\right), \text{ for } t \in [-\zeta, \zeta], \quad (4)$$

with σ^2 the variance of the Gaussian, $[-\zeta, \zeta]$ the support of the truncated Gaussian, and Υ the normalizing constant. Further, based on studies in [1, 5], we model the DCT coefficients using a zero-mean Laplacian distribution.¹

$$P(X = t) = \frac{\lambda}{2} \exp(-\lambda|t|), \quad (5)$$

¹DC components are typically modeled using Gaussian distributions with non-zero mean. To avoid the errors associated with estimating the Gaussian's mean and for simplicity, we assume that the DC coefficient can also be modeled using a zero-mean Laplacian distribution with zero mean.

We have assumed that the parameter λ is known; in practice, we estimate λ from the previously compressed image for each DCT frequency as described later in this section. From (5), we have

$$\begin{aligned} L_\lambda(kq) &:= P(\bar{X}_q = kq \mid q, k \in \mathbb{Z}) \\ &= \int_{(k-0.5)q}^{(k+0.5)q} \frac{\lambda}{2} \exp(-\lambda|\tau|) d\tau. \end{aligned} \quad (6)$$

Hence

$$P(\bar{X}_q = t \mid q) = \sum_{k \in \mathbb{Z}} \delta(t - kq) L_\lambda(kq). \quad (7)$$

Now, assuming that the round-off error Γ is independent of X and q , the distribution of \tilde{X} is obtained by convolving the distributions for \bar{X} and Γ (see Fig. 2). That is,

$$\begin{aligned} P(\tilde{X} = t \mid q) &= \int P(\bar{X}_q = \tau \mid q) P(\Gamma = t - \tau) d\tau \quad (8) \\ &= \begin{cases} \sum_{k \in \mathbb{Z}} \frac{1}{\Upsilon} \exp\left(-\frac{|t-kq|^2}{2\sigma^2}\right) L_\lambda(kq), \\ \quad \text{for } |t - kq| \in [-\zeta, \zeta], \\ 0, \quad \text{otherwise.} \end{cases} \end{aligned} \quad (9)$$

Given a set Ω of DCT coefficients at a particular frequency that are obtained from a previously compressed image, we can obtain the MAP estimate q^* of the quantization step used during previous compression assuming the DCT coefficients are independent as

$$q^* = \arg \max_{q \in \mathbb{Z}^+} P(\Omega, q) \quad (10)$$

$$= \arg \max_{q \in \mathbb{Z}^+} \left(\prod_{\tilde{X} \in \Omega} P(\tilde{X} \mid q) P(q) \right), \quad (11)$$

where $P(q)$ denotes the prior on q . Hence we can employ the following estimation algorithm:

1. Compute set of the desired frequency DCT coefficients Ω from the previously compressed image.
2. Estimate the parameter λ from the observations as

$$\lambda = \frac{N}{\sum_{\tilde{X} \in \Omega} |\tilde{X}|},$$

where N denotes the number of coefficients in the set Ω .

3. Assuming all quantization step-sizes are equally likely, use (9) with suitable parameters σ^2 and ζ to estimate

$$q^* = \arg \max_{q \in \mathbb{Z}^+} \left(\prod_{\tilde{X} \in \Omega} P(\tilde{X} \mid q) \right). \quad (12)$$




Fig. 2. Histogram of quantized DCT coefficients. The DCT coefficients from DCT frequency (4,4) of the grayscale Lenna image were subjected to quantization with step-size $q = 10$ during JPEG compression and then decompressed. Due to roundoff errors, the DCT coefficients are perturbed from integer multiples of 10.

The above algorithm is a refinement of the technique proposed by Fan and de Queiroz in [2]. While the core ideas remain the same, the final derived equation (9) differs because of significant variations in the starting points for the derivation and in the intermediate assumptions. Further, our derivation explicitly accounts for all the normalization constants, thereby allowing us to extend the above approach to estimate the CH of color images.

4. CH ESTIMATION FOR COLOR IMAGES

We build on the quantization step-size estimation algorithm for grayscale images to estimate the CH of color images. For color images, in addition to quantization, JPEG performs color transformation and subsampling along with the complementary interpolation. We realize that the DCT coefficient histograms of each color plane exhibit the near-periodic structure of Fig. 2 introduced by quantization only when the image is transformed to the original compression color space and the all interpolation effects are inverted. Otherwise, the near-periodic structure is not visible. This realization enables us to obtain the MAP estimate of a color image’s CH as in (1) by extending the statistical framework for grayscale images a straightforward way. Let $\Omega_{G,S}$ denote the set of DCT coefficients $\tilde{X}_{G,S}$ obtained by first transforming the image from F to the G color space representation, then undoing the interpolation S , and finally taking the DCT of the color planes. Then,

$$\begin{aligned} \{G^*, S^*, Q^*\} &= \arg \max_{G,S,Q} P(\text{Image}|G, S, Q)P(G, S, Q), \\ &= \arg \max_{G,S,Q} P(\Omega_{G,S}|G, S, Q)P(G)P(S)P(Q), \\ &= \arg \max_{G,S,Q} \prod_{\tilde{X}_{G,S} \in \Omega_{G,S}} P(\tilde{X}_{G,S}|G, S, Q)P(G)P(S)P(Q), \end{aligned}$$

assuming that the G , S , and Q are independent. The conditional probability $P(\tilde{X}_{G,S}|G, S, Q)$ of the DCT coefficients is set to (9), which is a metric for how well the image DCT coefficients conform to a near-periodic structure. If G , S , and Q were employed during the previous JPEG compression, then the histogram of $\Omega_{G,S}$ would be nearly periodic. Consequently, the associated $P(\tilde{X}_{G,S}|G, S, Q)$ would be large, and the MAP estimate would be accurate.

Ideally, the MAP estimation would require a search over all G and S . For practical considerations, we constrain our search to a dictionary of commonly employed compression color spaces and interpolations. The steps of our simple dictionary-based CH estimation algorithm are as follows:

1. Transform the observed color image to a test color space G .

2. Undo the effects of the test interpolation S . To undo interpolation by simple repetition, simply downsample the color plane. To undo interpolation by repetition and smoothing, first deconvolve the smoothing using a simple Tikhonov-regularized deconvolution filter [6], and then downsample the color plane.
3. Employ the quantization step-size estimation algorithm of Section 3 on the coefficients at each DCT frequency and each color plane.
4. Output the color transform and interpolation yielding the maximum conditional probability along with the associated quantization tables from Step 3.

The computational complexity of the dictionary-based CH estimation algorithm is determined by the size of the image, the number of the test color spaces, and the number of test subsamplings and interpolations. We can easily prune the number of test color spaces and interpolations to reduce the computational complexity by using a small part of the image. The quantization table estimates can be perfected using the entire image after the color space and interpolation is quickly determined.

5. RESULTS

We demonstrate via simulations that our algorithm can precisely estimate the CH of a previously JPEG-compressed color image. Further, we also demonstrate that recompression results obtained by exploiting the estimated CH are superior to results obtained by naively recompressing a color image using arbitrary compression settings.

To test our CH estimation algorithm, we chose a 480×480 -pixel uncompressed color image. We JPEG-compressed the image in the 8-bit CIE Lab color space using the sRGB to 8-bit CIE Lab color transformation [7] and employed $2 \times 2, 1 \times 1, 1 \times 1$ subsampling; that is, the luminance L color plane is not downsampled, while the chrominance planes a and b are downsampled by a factor of 2 in the horizontal and vertical directions. We employed the quantization tables described in [1, p. 37]. During decompression, the a and b planes are interpolated by first upsampling using repetition and then smoothing in the horizontal and vertical directions using a $\frac{1}{4} \times [1 \ 2 \ 1]$ impulse response filter.

For CH estimation, we test all color transforms from a dictionary consisting of RGB to YCbCr, Computer RGB to ITU.BT-601 YCbCr, Studio RGB to ITU.BT-601 YCbCr, RGB to Kodak PhotoYCC, sRGB to Linear RGB, sRGB to 8-bit CIE Lab, and sRGB to CMY transforms [7]. For each transform, we consider subsampling factors $2 \times 2, 1 \times 1, 1 \times 1$ (with and without smoothing during interpolation) and $1 \times 1, 1 \times 1, 1 \times 1$. As described in our algorithm in Section 4, we estimate the quantization table for each color transform and interpolation. By comparing the corresponding probabilities, we precisely identified that the sRGB to 8-bit CIE Lab color transformation was employed with $2 \times 2, 1 \times 1, 1 \times 1$ subsampling during the previous compression, and smoothing was employed during the decompression.

Our quantization step-size estimates, especially at the more important low DCT frequencies, is also accurate; the estimation errors (actual – estimated) for the L , a , and b color planes respectively are shown in Fig. 3. The \times indicates that the quantization step-size estimation for the corresponding DCT frequency was not possible because all the DCT coefficients were quantized to zero during the compression. The estimates for the a and the b planes suffer from seemingly large errors. For example, we incur an error of 92 while estimating an entry in the a plane’s quantization table; the actual

0	0	0	0	0	0	1	×	×	0	0	1	2	×	92	×	×	0	0	0	2	×	90	90	×
0	0	0	0	0	1	×	×	×	0	1	2	×	×	×	×	×	0	1	2	×	×	×	×	×
0	0	0	0	0	1	2	×	×	1	2	×	×	×	×	×	×	1	2	×	×	×	×	×	×
0	0	0	0	1	×	×	×	×	4	×	×	×	×	×	×	×	4	10	×	×	×	×	×	×
0	0	0	1	3	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
0	0	1	3	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
1	2	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×

Fig. 3. Quantization table estimation errors (actual – estimated) for the L, a and b color planes.

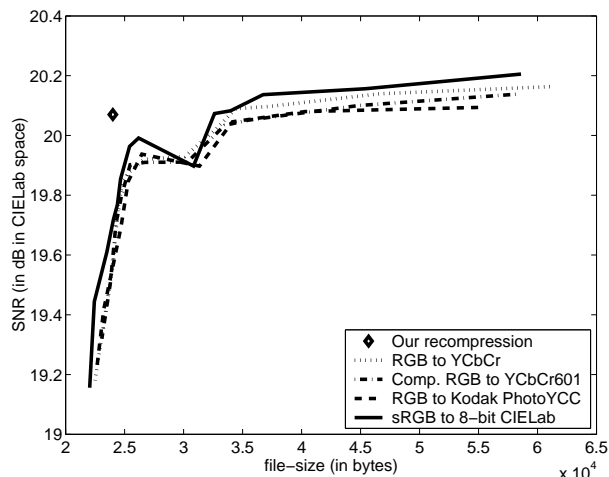


Fig. 4. File-size versus SNR tradeoff. The plot demonstrates that JPEG recompression using the estimated CH strikes an excellent file-size versus SNR tradeoff. Recompression using other compression settings suffer significant decreases in SNR to obtain even marginally lower file-sizes than our recompression.

quantization step-size was 100 but our algorithm’s estimate was 8. The error is a result of additional noise (compared to the no subsampling case) introduced by the necessary deconvolution step. However, the estimation error does not adversely affect applications such as recompression because, in reality, all the DCT coefficients at the corresponding frequency were set to zero during quantization.

The CH information enables us to effectively recompress the previously compressed image (see Fig. 4). To perform the recompression, we first transform the image into the compression color space using the estimated color transformation sRGB to 8-bit CIE Lab. Then, we deconvolve the effect of the smoothing employed during previous decompression on the a and b color planes. We then employ $2 \times 2, 1 \times 1, 1 \times 1$ subsampling and the estimated quantization tables (with 100 instead of the \times entries) to JPEG-compress the 8-bit CIE Lab color planes of the image using the popular IJG JPEG implementation [4]. Our recompression yielded a file-size of 24×10^3 bytes with an SNR of 20.07 dB; the SNR is computed in dB with respect to the original image in the perceptually-uniform CIE Lab color space.

For comparison, we also recompress the image using the RGB to YCbCr, Computer RGB to ITU.BT-601 YCbCr, RGB to Kodak PhotoYCC, and sRGB to 8-bit CIE Lab color transforms with $2 \times 2, 1 \times 1, 1 \times 1$ subsampling. We set the quantizer tables by setting the quality factor (QF) in the IJG JPEG implementation. The

QF is a reference number between 1 to 100; QF=100 set all the quantizer steps are unity and thus yields the best quality JPEG can possibly achieve. For different QF values, we noted the file-size and SNR (in dB in the CIE Lab space) pairs for all color transforms. Figure 4 compares the file-size versus SNR tradeoff achieved by the CH-powered recompression with the file-size versus SNR curves for naive recompression in the four color spaces. The plot confirms that exploiting the estimated CH information enables us to strike an excellent file-size versus distortion trade-off. Recompression results obtained using other JPEG compression settings either have to suffer a significant decrease in SNR to obtain even a slightly lower file-size or have to endure a significant increase in file-size to obtain marginal improvements in SNR. To attain a file-size of 24×10^3 bytes, recompression using the RGB to YCbCr, Computer RGB to ITU.BT-601 YCbCr, RGB to Kodak PhotoYCC, and sRGB to 8-bit CIE Lab transforms would have to tolerate 19.50 dB, 19.6 dB, 19.64 dB and 19.72 dB SNRs respectively (compared to 20.07 dB SNR for our recompression).

6. CONCLUSIONS

We have proposed a new statistical framework to estimate a color image’s JPEG compression history (CH) by exploiting the near-periodic structure of discrete cosine transform coefficients created by the previous JPEG’s quantization step. Our algorithm exhibits good estimation performance. Further, JPEG recompression performed using the estimated CH strikes an excellent file-size versus distortion tradeoff.

7. REFERENCES

- [1] W. Pennebaker and J. Mitchell, *JPEG, Still Image Data Compression Standard*, Van Nostrand Reinhold, 1993.
- [2] Z. Fan and R. de Queiroz, “Identification of bitmap compression history: JPEG detection and quantizer estimation,” *IEEE Trans. Image Processing*, vol. 12, pp. 230–235, Feb. 2003.
- [3] R. Neelamani, R. de Queiroz, and R. G. Baraniuk, “Compression color space estimation of JPEG images using lattice basis reduction,” in *Proc. IEEE ICIP ’01*, Thessaloniki, Greece, Oct. 2001, pp. 890–893.
- [4] *Independent JPEG Group Library*, www.ijg.org.
- [5] R. J. Clarke, *Transform Coding of Images*, Academic Press, London, England, 1985.
- [6] A. N. Tikhonov and V. Y. Arsenin, *Solutions of Ill-Posed Problems*, V. H. Winston & Sons, Washington D.C., 1977.
- [7] Charles Poynton, *A Technical Introduction to Digital Video*, Wiley, New York, 1996.