JPEG Compression History Estimation for Color Images

Ramesh Neelamani[*], *Ricardo de Queiroz, Zhigang Fan, Sanjeeb Dash, and Richard G. Baraniuk*

Accepted by the IEEE Transactions on Image Processing

Abstract

We routinely encounter digital color images that were previously JPEG-compressed. En-route to the image's current representation, the previous JPEG compression's various settings—termed its JPEG compression history (CH)–are often discarded after the JPEG decompression step. Given a JPEG-decompressed color image, this paper aims to estimate its lost JPEG CH. We observe that the previous JPEG compression's quantization step introduces a lattice structure in the discrete cosine transform (DCT) domain. This paper proposes two approaches that exploit this structure to solve the JPEG Compression History Estimation (CHEst) problem. First, we design a statistical dictionary-based CHEst algorithm that tests the various CHs in a dictionary and selects the *maximum a posteriori* estimate. Second, for cases where the DCT coefficients closely conform to a 3-D parallelepiped lattice, we design a *blind* lattice-based CHEst algorithm. This algorithm exploits the fact that the JPEG CH is encoded in the nearly orthogonal bases for the 3-D lattice and employs novel lattice algorithms and recent results on nearly orthogonal lattice bases to estimate the CH. Both algorithms provide robust JPEG CHEst performance in practice. Simulations demonstrate that JPEG CHEst can be extremely useful in JPEG recompression; the estimated CH allows us to recompress a JPEG-decompressed image with minimal distortion (large signal-to-noise-ratio) and simultaneously achieve a small file-size.

Keywords: JPEG, compression, color, history, recompression, lattice, quantization

[*]R. Neelamani is with the ExxonMobil Upstream Research Company, 3319 Mercer, Houston, TX 77027–6019. Email: neelsh@rice.edu. Fax: 713 431 6161. R. de Queiroz is with the Department of Electrical Engineering, Universidade de Brasilia, CP 04591, Brasilia, DF, 70910-900, Brazil. Email: queiroz@ieee.org. Z. Fan is with Xerox Research and Technology, Xerox Corporation, 800 Phillips Road, Webster, NY 14580, USA. Email: ZFan@crt.xerox.com. S. Dash is with the IBM T. J. Watson Research Center, Yorktown Heights, NY 10598, USA. Email: sanjeebd@us.ibm.com. R. G. Baraniuk is with the Department of Electrical and Computer Engineering, Rice University, 6100 South Main Street, Houston, TX 77005–1892, USA. Email: richb@rice.edu. R. Neelamani and R. G. Baraniuk were both supported by grants from NSF, AFOSR, ONR, DARPA, and Texas Instruments Leadership University Program. Web: www.dsp.rice.edu. Contact author: R. Neelamani.

1 Introduction

A digital color image is a collection of pixels with each pixel a 3-dimensional (3-D) color vector. The vector elements specify the pixel's color with respect to a chosen color space; for example, *RGB*, *YCbCr*, et cetera [1, 2]. Joint Photographic Experts Group (JPEG) is a commonly used standard to compress digital color images [3]. JPEG compresses by quantizing the discrete cosine transform (DCT) coefficients of the image's three color planes; see Fig. 1 for an overview. However, the various settings employed during JPEG compression and decompression are not standardized [3]. The following JPEG settings can be chosen by the user or an imaging device: (i) the color space used to compress the image's three color planes independently; (ii) the subsampling employed on each color plane during compression and the complementary interpolation employed during decompression; and (iii) the quantization table used to compress each color plane. We refer to these settings 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 often imported into Microsoft Powerpoint or Word documents using graphics programs such as Microsoft Clip Gallery and then stored internally using a decompressed format. JPEG images are also routinely converted to lossless-compression formats such as Windows bitmap (BMP) format (say, to create a background image for Windows or to feed a print driver) or Tagged Image File Format (TIFF). In such cases, the JPEG compression settings are discarded after decompression.

We aim to estimate the JPEG CH from a given JPEG-decompressed color image. We refer to this problem as JPEG Compression History Estimation (JPEG CHEst).

The CH, if available, can be used for a variety of applications. The file-size of a JPEG image is typically *significantly* smaller than the file-size after the image's conversion to BMP or TIFF format. The JPEG CH enables us to effectively recompress such converted BMP and TIFF images; JPEG-compressing the image with previous JPEG settings yields significant file-size reduction without introducing additional distortion. The JPEG CH can also be used by "smart" print servers to reduce artifacts from received BMP images such as blocking due to previous JPEG compression. To alleviate such artifacts by adapting techniques described in [4, 5], the print server would need the image's JPEG CH. 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.

The CHEst problem is relatively unexplored. Fan and de Queiroz have proposed a statistical framework to perform CHEst for gray-scale images [6]; for a gray-scale image, the CH comprises only the quantization table employed during previous JPEG operations. However, CHEst for color images remains unexplored.

This paper proposes two new frameworks to perform CHEst for color images.

First, we derive a statistical framework for CHEst. We observe that JPEG leaves its signature by quantizing the image's DCT coefficients and forcing them to conform to near-periodic structures. We statistically characterize this near-periodicity for a single color plane. The resulting framework can be exploited to estimate a gray-scale image's CH, namely, its quantization table. We extend the statistical framework to color images and design a dictionary-based CHEst approach. The dictionary consists of typical color transformations, subsampling factors, and interpolations. We adopt a *maximum a posteriori* (MAP) approach to estimate the color image's CH from the dictionary

$$\left\{\widehat{G}, \widehat{S}, \widehat{Q}\right\} = \arg\max_{G,S,Q} P(\text{Image}, G, S, Q)\,, \tag{1}$$

with $P(\cdot)$ denoting the probability and $\widehat{G}$, $\widehat{S}$, $\widehat{Q}$ the estimated compression color space, subsampling and associated interpolation, and quantization tables, respectively.

Second, we consider the case when the transform from the color space used to perform quantization to the image's current representation color space is affine and when no subsampling is employed during JPEG compression. For such a case, we develop a novel, *blind*, lattice-based CHEst algorithm. Such a blind approach is required when an unknown proprietary color transform is employed by JPEG. We realize that, after JPEG decompression, such an image's DCT coefficients closely conform to a 3-D parallelepiped[1] *lattice* structure determined by the affine color transform. Formally, a *lattice* is a set of integer linear combinations of a given set of vectors. The minimal set of vectors whose integer linear combinations span all lattice points is a *lattice basis*. We also realize that the JPEG CH information is encoded in the nearly orthogonal bases that span the DCT lattices. Recently, [7, 8] derived the geometric conditions for a lattice basis to contain the shortest non-zero lattice vector and the conditions to characterize the such bases' uniqueness. Using these recent insights, and using novel applications of existing lattice algorithms, we estimate the color image's CH, namely, the affine color transform and the quantization tables.

The proposed CHEst algorithms demonstrate excellent performance in practice. Further, we verify that CHEst allows us to recompress an image with minimal distortion (large signal-to-noise-ratio (SNR)) and simultaneously achieve a small file-size (see Figs. 7 and 8).

The rest of this paper is organized as follows. We first provide a brief overview of color transforms and JPEG in Sections 2 and 3. We derive the statistical framework for CHEst for gray-scale images in Section 4 and extend this framework to design dictionary-based CHEst for color images in Section 5. In

[1] A solid with six faces, each of which is a parallelogram.

Section 6, we describe the 3-D lattice structure of a JPEG-decompressed image when JPEG uses an affine color transform and no subsampling. Section 7 overviews the properties of nearly orthogonal lattice bases and some celebrated CHEst-relevant lattice problems. In Section 8, we describe lattice-based CHEst and its experimental performance. We demonstrate CHEst's utility in JPEG recompression in Section 9 and conclude in Section 10.

2 Color Spaces and Transforms

Color perception is a sensation produced when light excites the receptors in the human retina. Color can be described by specifying the light's spectral power distribution. Such a description is highly redundant because the human retina has only three types of receptors that influence color perception.[2] Consequently, three numerical components are sufficient to describe a color; this is termed the *trichromatic theory* [2].

Based on the trichromatic theory, digital color imaging devices use three parameters to specify any color; the three parameters can be viewed as a 3-D vector. The *color space* is the reference coordinate system with respect to which the 3-D vector describes color [1, 2]. There exist many different coordinate systems or color spaces according to which a color can be specified. For example, the Commission Internationale de L'Éclairage (CIE) defined the *CIE XYZ* color space to specify all visible colors using positive *X*, *Y*, and *Z* values [1, 2]. Other examples include different varieties of *RGB* (Red *R*, Green *G*, and Blue *B*) and *YCbCr* (luminance *Y*, and chrominances *Cb* and *Cr*) color spaces. These color spaces are related to each other and to reference color spaces such as the *CIE XYZ* via linear or non-linear color transformations. For example, the popular Independent JPEG Group (IJG) JPEG implementation [9] converts the digital color image's $0-255$-valued *R*, *G*, *B* components to $0-255$-valued *Y*, *Cb*, *Cr* components using the following transformation

$$\begin{bmatrix} Y \\ Cb \\ Cr \end{bmatrix} = \begin{bmatrix} 0.299 & 0.587 & 0.114 \\ -0.169 & -0.331 & 0.5 \\ 0.5 & -0.419 & -0.081 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix} + \begin{bmatrix} 0 \\ 128 \\ 128 \end{bmatrix}. \tag{2}$$

The resulting *YCbCr* space is also referred to as the *ITU.BT-601 YCbCr* space [1]. The inverse color trans-

[2] A fourth type of receptor is also present in the retina, but it does not affect color perception because it is effective only at extremely low light levels [2].

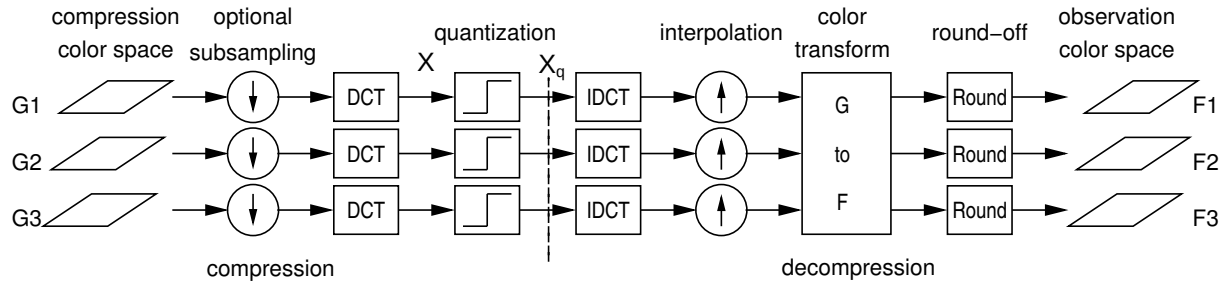

Figure 1: *Overview of JPEG compression and decompression.*

formation from the *ITU.BT-601 YCbCr* space to the *RGB* space is given by

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} = \begin{bmatrix} 1.0 & 0.0 & 1.402 \\ 1.0 & -0.344 & -0.714 \\ 1.0 & 1.772 & 0.0 \end{bmatrix} \left(\begin{bmatrix} Y \\ Cb \\ Cr \end{bmatrix} - \begin{bmatrix} 0 \\ 128 \\ 128 \end{bmatrix} \right). \tag{3}$$

The transforms in both (2) and (3) are *affine*. Henceforth, we refer to the 3×3 matrix as the affine transform's *linear component* and the 3×1 shift as the affine transform's *additive component*.

Later in this chapter, we will invoke a variety of color spaces that are inter-related by affine or non-linear transforms. We refer the reader to [1, 2] for additional information on color, different color spaces, and transforms.

3 Effects of JPEG Compression and Decompression

In this section, we review the CHEst-relevant JPEG compression and decompression steps. We do not describe all JPEG operations, but present a model that fully accounts for the effects of JPEG compression and decompression on an image. The model folds quantization and de-quantization into one step, and ignores all entropy coding steps because these do not affect the final image. For further JPEG details, we refer the reader to [3].

Consider an observed color image represented in the hypothetical 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 the *compression color space*.

JPEG compression essentially 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 non-overlapping 8×8 blocks. Take the 2-D DCT of each block.

3. Quantize the coefficients at each DCT 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 the quantization step-size for the corresponding DCT frequency, then the quantized DCT coefficient $\overline{X}_q$ is obtained by

$$\overline{X}_q := \mathrm{round}\left(\frac{X}{q}\right) q. \tag{4}$$

   See Fig. 2 for examples of quantization tables; each entry in the 8×8 quantization table is the quantization step-size for an 8×8 image block's corresponding DCT coefficient.

JPEG decompression performs the following operations:

1. Compute 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 [9] 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]

Henceforth, we will refer to the zero frequency DCT coefficient as the *DC* coefficient and the remaining 63 DCT coefficients as the *AC* coefficients.

4 CHEst for Gray-Scale Images

For gray-scale images, JPEG compression and decompression replicates the steps outlined in Section 3 for a single color plane but without subsampling and interpolation. Hence, the CHEst problem simplifies to estimating the quantization tables employed during the previous JPEG compression. Due to JPEG's quantization operations, the JPEG-decompressed gray-scale images' DCT coefficient histograms exhibit a

[3] In reality, round-offs occur also after the inverse DCT step. The model consolidates all the round-off operations into one step for the sake of simplicity.

10	7	6	10	14	24	31	37
7	7	8	11	16	35	36	33
8	8	10	14	24	34	41	34
8	10	13	17	31	52	48	37
11	13	22	34	41	65	62	46
14	21	33	38	49	62	68	55
29	38	47	52	62	73	72	61
43	55	57	59	67	60	62	59

Quantization table 1

10	11	14	28	59	59	59	59
11	13	16	40	59	59	59	59
14	16	34	59	59	59	59	59
28	40	59	59	59	59	59	59
59	59	59	59	59	59	59	59
59	59	59	59	59	59	59	59
59	59	59	59	59	59	59	59
59	59	59	59	59	59	59	59

Quantization table 2

Figure 2: *Examples of JPEG quantization tables for 8×8 DCT blocks.*

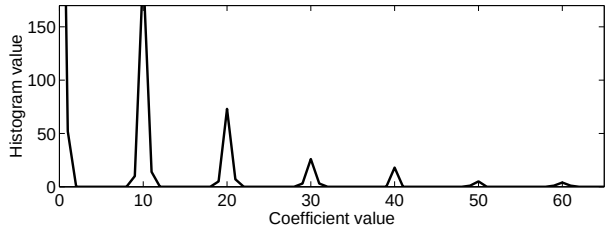

Figure 3: *Histogram of quantized DCT coefficients. The DCT coefficients from DCT frequency (4,4) of the gray-scale Lena image were subjected to quantization with step-size* $q = 10$ *during JPEG compression and then decompressed. Due to roundoff errors, the DCT coefficients are perturbed from integer multiples of* 10.

near-periodic structure with the period determined by the quantization step-size. In this section, to estimate the quantization table, we derive a statistical framework that characterizes the near-periodic structure.

4.1 Statistical framework

An arbitrary DCT coefficient $\widetilde{X}$ of a JPEG-decompressed gray-scale image can be obtained by adding to the corresponding quantized coefficient $\overline{X}_q$ (see (4)) a round-off error term Γ

$$\widetilde{X} = \overline{X}_q + \Gamma. \tag{5}$$

As described in [6], we can model Γ using a truncated Gaussian distribution

$$P(\Gamma = t) = \Upsilon \exp\left(-\frac{t^2}{2\sigma^2}\right), \qquad \text{for } t \in [-\zeta, \zeta], \tag{6}$$

with σ^2 the Gaussian's variance, $[-\zeta, \zeta]$ the truncated Gaussian's support, and Υ the normalizing constant. (For example, $\sigma^2 = 0.8$ and $\zeta = 6$.) Further, based on studies in [3, 10], we can model the DCT coefficients using a zero-mean Laplacian distribution

$$P(X = t) = \frac{\lambda}{2} \exp\left(-\lambda|t|\right). \tag{7}$$

We have assumed that the parameter λ is known; in practice, we estimate λ from the observed decompressed image for each DCT frequency as described later in this section. From (7), we have

$$L_\lambda(kq) := P\big(\overline{X}_q = kq \mid q,\ k \in \mathbb{Z}\big) = \int_{(k-0.5)q}^{(k+0.5)q} \frac{\lambda}{2} \exp\left(-\lambda|\tau|\right)\, d\tau \tag{8}$$

and hence

$$P\big(\overline{X}_q = t \mid q\big) = \sum_{k \in \mathbb{Z}} \delta\left(t - kq\right) L_\lambda(kq). \tag{9}$$

Now, assuming that the round-off error Γ is independent of X and q, $\widetilde{X}$'s distribution is obtained by convolving the distributions for $\overline{X}$ and Γ (see Fig. 3). That is,

$$P\Big(\widetilde{X} = t \mid q\Big) = \int P\big(\overline{X}_q = \tau \mid q\big)\, P(\Gamma = t - \tau)\, d\tau \tag{10}$$

$$= \begin{cases} \sum_{k \in \mathbb{Z}}\ \Upsilon \exp\left(-\frac{|t-kq|^2}{2\sigma^2}\right) L_\lambda(kq), & \text{for } |t - kq| \in [-\zeta, \zeta], \\ 0, & \text{otherwise.} \end{cases} \tag{11}$$

Let $\mathcal{D}_i$ denote the set of the JPEG-decompressed image's i^{th}-frequency DCT coefficients; that is, $\mathcal{D}_i$ comprises one coefficient from each 8×8 block. Given $\mathcal{D}_i$, we can obtain the MAP estimate $\widehat{q}_i$ of the quantization step used on the i^{th}-frequency coefficients during previous compression as

$$\widehat{q}_i = \arg\max_{q \in \mathbb{Z}^+} P(\mathcal{D}_i, q) \tag{12}$$

$$= \arg\max_{q \in \mathbb{Z}^+} \left(\prod_{\widetilde{X} \in \mathcal{D}_i} P\Big(\widetilde{X} \mid q\Big)\, P(q) \right), \tag{13}$$

where the DCT coefficients are assumed to be independent and $P(q)$ denotes the prior on q.

4.2 Algorithm steps

Using the statistical framework derived in the previous section, we can estimate the 8×8 quantization table $Q := \{q_i\}$, with $i = 1, \ldots, 64$, enumerating the 64 DCT frequencies, using the following steps:

1. For each frequency i, compute the set $\mathcal{D}_i$ of the observed decompressed image's DCT coefficients.

2. Estimate the parameter λ from the observations as

$$\lambda = \frac{N}{\sum_{\widetilde{X}\in\mathcal{D}_i} |\widetilde{X}|},$$

   with N the number of coefficients in the set $\mathcal{D}_i$.

3. Assuming a uniform prior on q, use (11) with suitable parameters σ^2 and ζ to estimate

$$\widehat{q}_i \;=\; \arg\max_{q\in\mathbb{Z}^+} \left(\prod_{\widetilde{X}\in\mathcal{D}_i} P\Big(\widetilde{X} \mid q\Big) \right). \tag{14}$$

This algorithm is not entirely new; it is a refinement of the technique proposed by Fan and de Queiroz in [6]. While the core ideas remain the same, the final derived equation (11) differs because of significant variations in the starting points for the derivation and in the intermediate assumptions. Further, our derivation explicitly accounts for all normalization constants, thereby allowing us to extend the above approach to estimate the CH of color images.

5 Dictionary-based CHEst for Color Images

In this section, we build on the quantization step-size estimation algorithm for gray-scale images from Section 4 to perform CHEst for color images.

5.1 Statistical framework

For color images, in addition to quantization, JPEG performs color transformation and subsampling along with the complementary interpolation. We observe that the DCT coefficient histogram of each color plane exhibits the near-periodic structure of Fig. 3 introduced by quantization only when the image is transformed to the original compression color space and all interpolation artifacts are removed. Hence we can obtain

the MAP estimate of a color image's CH as in (1) using a simple extension of the statistical framework for gray-scale images.

Let $\widetilde{X}_{Gj,S,i}$ denote the set of i^{th}-frequency DCT coefficients from the j^{th}-color plane, $j = 1, 2, 3$. Assume that $\widetilde{X}_{Gj,S,i}$ is obtained by first transforming the image from the F to the G color space representation, then undoing the interpolation S, and finally taking the DCT of the color planes. Let $\mathcal{D}_{Gj,S,i}$ denote the set of all $\widetilde{X}_{Gj,S,i}$. Then,

$$\begin{aligned}
\Big\{\widehat{G}, \widehat{S}, \widehat{Q}\Big\} &= \arg\max_{G,S,Q} P(\text{Image}|G,S,Q)\, P(G,S,Q) \\
&= \arg\max_{G,S,Q} \prod_{i=1}^{64} \prod_{j=1}^{3} P(\mathcal{D}_{Gj,S,i}|G,S,Q)\, P(G)\, P(S)\, P(Q) \\
&= \arg\max_{G,S,Q} \prod_{i=1}^{64} \prod_{j=1}^{3} \prod_{\widetilde{X}_{G,S} \in \mathcal{D}_{G,S}} P\Big(\widetilde{X}_{Gj,S,i}|G,S,Q\Big)\; P(G)\, P(S)\, P(Q), \qquad (15)
\end{aligned}$$

assuming that all the $\widetilde{X}_{Gj,S,i}$ coefficients and the choices of G, S, and Q are independent. In (15), the DCT coefficients' conditional probability $P(\widetilde{X}_{Gj,S,i}|G,S,Q)$ is computed using (11), which is a metric for how well the image DCT coefficients conform to a near-periodic structure. Hence, if G, S, and Q were actually employed during the previous JPEG compression, then the histogram of the i^{th}-frequency DCT coefficients would be nearly periodic and the associated $P(\widetilde{X}_{Gj,S,i}|G,S,Q)$ would be large. Consequently, the MAP estimate would be accurate.

5.2 Algorithm steps

In general, the MAP estimation in (15) would require a search over all G and S. For practical considerations, we constrain the search to a dictionary comprising commonly employed compression color spaces and interpolations. Dictionary-based CHEst steps are as follows:

1. Choose a test color space G and interpolation method S from the dictionary.
2. Transform the observed color image to the color space G.
3. 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 [11] and then downsample the color plane.

4. Employ the quantization table estimation step from Section 4 on each color plane.

5. Try all the G and S in the dictionary, and output the G and S yielding the maximum conditional probability (15) along with the associated quantization tables from Step 4.

Dictionary-based CHEst's computational complexity is determined by the image size, the number of the test color spaces, and the number of test subsamplings and interpolations in the dictionary. In practice, we can easily and reliably eliminate a majority of the dictionary elements using just a small part of the image. The CH can then be estimated quickly by applying the dictionary-based CHEst with the pruned dictionary on the entire image.

5.3 Dictionary-based CHEst results

Dictionary-based CHEst precisely estimates a JPEG-decompressed color image's CH when the dictionary contains the actual color transform and interpolation. We demonstrate dictionary-based CHEst's performance using the 512×512 *Lena* color image [12] and a specific JPEG CH choice. The algorithm performed equally well on a wide variety of experiments comprising different images and compression color spaces. Our Matlab scripts can be downloaded from *www.dsp.rice.edu/software*. We JPEG-compressed *Lena* in the *8-bit CIELab* color space using the *sRGB* to 8-bit CIELab color transformation [1] and employed $2 \times 2, 1 \times 1, 1 \times 1$ subsampling; that is, the luminance *L* color plane was not downsampled, while the chrominance planes *a* and *b* were downsampled by a factor of 2 in the horizontal and vertical directions. We employed quantization tables 1 from Fig. 2 for the *L* plane and quantization table 2 from Fig. 2 for the both the *a* and *b* planes. During decompression, the *a* and *b* planes were 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. This decompressed image is the input to dictionary-based CHEst.

To perform CHEst, we tested all color transforms from a dictionary consisting of *RGB* to *ITU.BT-601 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 CIELab*, and *sRGB* to *CMY* transforms [1]. For each transform, we considered subsampling factors $2 \times 2, 1 \times 1, 1 \times 1$ (with and without smoothing during interpolation) and $1 \times 1, 1 \times 1, 1 \times 1$.

During the conditional probability computations (15), we assumed that all color transforms and quantization step-sizes are equally likely; that is, set $P(G) = P(Q) = 1$. When larger subsampling factors and smoothing are employed, the DCT coefficients deviate further from their quantized values, resulting in relatively lower conditional probabilities. To level this effect, we need to adapt σ^2 and the priors $P(S)$. To

test if a color plane was subsampled by a factor of 2 and then smoothed during interpolation, we set the $\sigma^2 = 0.8$ (see (11)) during the quantization table estimation step. To test if no smoothing was employed during interpolation, we set $\sigma^2 = 0.75$, and to test if no subsampling was employed, we reduced the σ^2 to 0.5. Further, we set the prior $P(S) = 0.55$ for the $2 \times 2, 1 \times 1, 1 \times 1$ with smoothing, $P(S) = 0.35$ for the $2 \times 2, 1 \times 1, 1 \times 1$ without smoothing, and $P(S) = 0.1$ for the $1 \times 1, 1 \times 1, 1 \times 1$ subsampling. We set $\zeta = 6$ in (11) during our experiments. These settings worked well on all our experimental tests.

By comparing the conditional probabilities' logarithms (listed in Table 1), we precisely identified that the *sRGB* to *8-bit CIELab* color transformation was employed with $2 \times 2, 1 \times 1, 1 \times 1$ subsampling during the previous compression, and that smoothing was employed during the decompression; the corresponding conditional probability value (enclosed by a $\square$ in Table 1) is the largest.

Figure 4 illustrates the algorithm's quantization table estimates (see Fig. 2 for the actual tables). Our quantization step-size estimates were quite accurate, especially at the more important low frequencies. Note that to estimate the quantization step-size, at least one coefficient should be quantized to a non-zero value. For many of the high frequencies, *all* coefficients were quantized to zero because the actual quantization step-sizes (see Fig. 2) were large. For such quantized-to-zero frequencies, our algorithm typically returned the maximum quantization step size included in our search range (100 in our case); they are marked by $\times$'s in Fig. 4. For some quantized-to-zero frequencies in the *a* and *b* planes, our algorithm did not return the maximum quantization step-size because the *a* and *b* plane coefficients are more noisy (compared to the *L* plane coefficients) due to the additional deconvolution step (Step 3 in Section 5.2). For example, our algorithm estimated one of the *a* plane's quantization step-sizes to be 7, whereas actual quantization step-size was 59. Note that such errors, though seemingly large, have negligible impact on applications such as recompression; any quantization step-size estimate greater than 1 would reset most of the DCT coefficients to zero (as desired) during recompression.

6 Blind CHEst and Lattices

The dictionary-based CHEst approach described in Section 5 would fail if an unknown proprietary color space was used to perform the JPEG compression. This motivates us to develop a *blind* approach that does not rely on a fixed dictionary of known color spaces. Blind lattice-based CHEst can handle cases where the transform from the compression color space to the current color space is affine and no subsampling is employed during the previous JPEG compression. *For such cases, Blind CHEst aims to estimate the affine transform and quantization tables employed during the previous JPEG compression from the JPEG-*

Table 1: *Logarithms of conditional probabilities ($\times 10^6$) for dictionary-based CHEst experiments.*

Color transform	$1 \times 1, 1 \times 1, 1 \times 1$	$2 \times 2, 1 \times 1, 1 \times 1$ (without smoothing)	$2 \times 2, 1 \times 1, 1 \times 1$ (with smoothing)
RGB to *ITU.BT-601 YCbCr*	-0.88	-0.88	-0.8
Computer RGB to *ITU.BT-601 YCbCr*	-0.83	-0.83	-0.75
Studio RGB to *ITU.BT-601 YCbCr*	-0.88	-0.89	-0.81
RGB to *Kodak PhotoYCC*	-0.79	-0.78	-0.69
sRGB to *Linear RGB*	-1.5	-1.9	-1.8
sRGB to *8-bit CIELab*	-0.73	-0.71	$\boxed{\mathbf{-0.53}}$
sRGB to *CMY*	-1.5	-1.9	-1.8

decompressed image. The lattice geometry of a JPEG decompressed image's DCT coefficients holds the key to blind JPEG CHEst.

6.1 Lattice fundamentals

Lattices are central to a number of fields including coding theory, number theory, and crystallography [13–16]. A lattice is the set of all linear integer combinations of a finite set of vectors. In $\mathbb{R}^n$, a lattice $\mathcal{L} := \{\mathcal{B}u : u \in \mathbb{Z}^m\}$, with $\mathcal{B}$ a real $n \times m$ matrix. Figures 5(a) and (b) are both illustrations of 3-D lattices. The columns of $\mathcal{B}$ are said to *span* the lattice $\mathcal{L}$. If $\mathcal{B}$ contains the minimal set of vectors spanning $\mathcal{L}$, then it is termed a *basis* for $\mathcal{L}$. A lattice can have more than one basis. Any two bases $\mathcal{B}_1$ and $\mathcal{B}_2$ for $\mathcal{L}$ have the same number of vectors and are related by $\mathcal{B}_1 = \mathcal{B}_2\mathcal{U}$, with $\mathcal{U}$ an *unimodular* matrix–an integer matrix with determinant equal to ± 1.

6.2 Ideal lattice structure of DCT coefficients

In the absence of round-off noise, due to JPEG's quantization step, a JPEG-decompressed color image' 3-D DCT vectors conform to a regular parallelepiped lattice structure.

Consider an arbitrary 8×8 uncompressed color image block in the G color space that the DCT acts on during JPEG compression. Let $X_{G1,i}$, $X_{G2,i}$, and $X_{G3,i}$ denote the respective i^{th}-frequency DCT coefficients of the $G1$, $G2$, and $G3$ planes in the chosen 8×8 color image block. As described in Section 3, JPEG

10	7	6	10	14	24	31	37
7	7	8	11	16	35	36	32
8	8	10	14	24	34	40	33
8	10	13	17	31	51	47	35
11	13	22	34	40	63	×	44
14	21	32	37	47	×	×	51
28	35	44	49	×	×	×	×
×	×	×	×	×	×	×	×

L's table

10	11	14	28	54	9	×	9
11	13	15	36	50	9	7	7
14	15	30	9	9	5	×	7
26	34	48	×	×	×	7	×
×	×	×	×	×	×	×	×
×	×	×	×	×	×	×	×
×	×	×	×	×	×	×	×
×	×	×	×	×	×	×	×

a's table

10	11	14	26	×	7	×	×
11	13	15	35	×	×	×	×
13	15	31	48	×	×	×	×
25	34	×	×	×	×	×	×
×	×	×	×	×	×	×	×
×	×	×	×	×	×	×	×
×	×	×	×	×	×	×	×
×	×	×	×	×	×	×	×

b's table

Figure 4: *Dictionary-based CHEst algorithm's quantization tables estimates for the* 8-bit CIELab's L, a, and b color planes.

quantizes each plane's DCT coefficients independently to

$$\overline{X}_{G,i} := \begin{bmatrix} \overline{X}_{G1,q_{i,1}} \\ \overline{X}_{G2,q_{i,2}} \\ \overline{X}_{G3,q_{i,3}} \end{bmatrix} := \begin{bmatrix} q_{i,1} & 0 & 0 \\ 0 & q_{i,2} & 0 \\ 0 & 0 & q_{i,3} \end{bmatrix} \begin{bmatrix} \text{round}\left(\frac{X_{G1,i}}{q_{i,1}}\right) \\ \text{round}\left(\frac{X_{G2,i}}{q_{i,2}}\right) \\ \text{round}\left(\frac{X_{G3,i}}{q_{i,3}}\right) \end{bmatrix} := \mathcal{Q}_i \begin{bmatrix} \text{round}\left(\frac{X_{G1,i}}{q_{i,1}}\right) \\ \text{round}\left(\frac{X_{G2,i}}{q_{i,2}}\right) \\ \text{round}\left(\frac{X_{G3,i}}{q_{i,3}}\right) \end{bmatrix}, \qquad (16)$$

with $q_{i,1}$, $q_{i,2}$, and $q_{i,3}$ the respective quantization step-sizes for the three color planes. Clearly, since $\left[\ \text{round}\left(\frac{X_{G1,i}}{q_{i,1}}\right),\ \text{round}\left(\frac{X_{G2,i}}{q_{i,2}}\right),\ \text{round}\left(\frac{X_{G3,i}}{q_{i,3}}\right)\right]^T \in \mathbb{Z}^3$ (the superscript T denotes matrix transpose), the vector $\left[\ \overline{X}_{q_{i,1}},\ \overline{X}_{q_{i,2}},\ \overline{X}_{q_{i,3}}\right]^T$ lies on a 3-D lattice with basis $\mathcal{Q}_i$. Figure 5(a) illustrates that the quantized 3-D DCT vector lies on a 3-D lattice; it is a *rectangular box* lattice because $\mathcal{Q}_i$'s columns are orthogonal to each other.

After quantization, assume that the image is subjected to an affine color transform from the G to the F color space. Let $\mathcal{T}$ denote the affine transform's linear component. Further assume that no round-off is

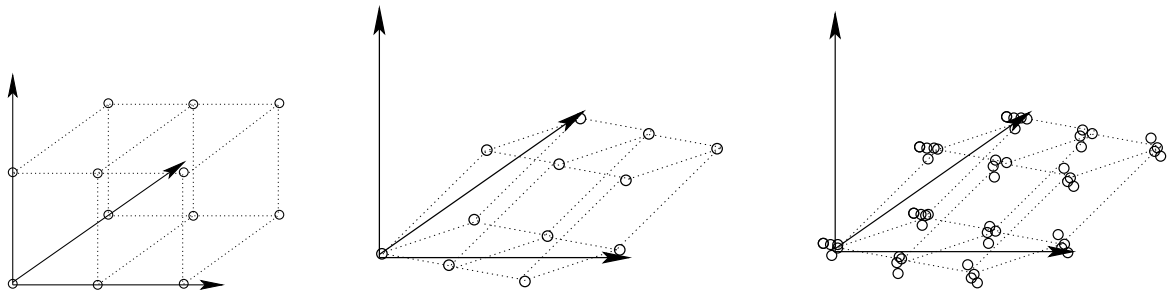

(a) G space (b) F space (no round-off) (c) F space (with round-off)

Figure 5: *Lattice structures in a JPEG-decompressed color image: (a) In the G space, all 3-D DCT vectors from the same DCT frequency lie on a rectangular lattice's vertices. The 3-D vectors are denoted by small circles. (b) Assuming round-off errors are absent, in the F space, the 3-D DCT vectors lie on a parallelepiped lattice's vertices. (c) Round-off errors slightly perturb the F space 3-D DCT vectors from the parallelepiped lattice locations.*

performed. Then, the transformed image's i^{th}-frequency ($i \neq$ DC) 3-D DCT vectors can be expressed as

$$X_{F,i} := \begin{bmatrix} X_{F1,i} \\ X_{F2,i} \\ X_{F3,i} \end{bmatrix} := \mathcal{T} \begin{bmatrix} \overline{X}_{G1,q_{i,1}} \\ \overline{X}_{G2,q_{i,2}} \\ \overline{X}_{G3,q_{i,3}} \end{bmatrix} = \mathcal{T}\mathcal{Q}_i \begin{bmatrix} \text{round}\left(\frac{X_{G1,i}}{q_{i,1}}\right) \\ \text{round}\left(\frac{X_{G2,i}}{q_{i,2}}\right) \\ \text{round}\left(\frac{X_{G3,i}}{q_{i,3}}\right) \end{bmatrix}, \tag{17}$$

Thus, in the F space representation, the i^{th}-frequency 3-D DCT vectors lie on a lattice $\mathcal{L}_i$ with basis $\mathcal{T}\mathcal{Q}_i$. (The affine transform's additive component affects only the DC coefficients; it shifts the DC coefficient lattice from the origin.) Figure 5(b) illustrates the lattice geometry of the 3-D DCT vectors in the F space; the vectors lie on a *parallelepiped* lattice because $\mathcal{T}\mathcal{Q}_i$'s columns are not orthogonal.

6.3 Round-offs perturb ideal lattice geometry

In reality, a JPEG-decompressed image is always subjected to round-off during the decompression (see Fig. 1 and Section 3). Hence, any 3-D DCT vector of the given JPEG-decompressed image can be expressed as

$$\widetilde{X}_{F,i} = X_{F,i} + N_i, \tag{18}$$

with N_i denoting the 3-D round-off noise vector. Based on (6), we can statistically model N_i as

$$P(\|N_i\|_2 = t) = \Upsilon \exp\left(-\frac{t^2}{2\sigma^2}\right), \qquad \text{for } t \in [-\zeta, \zeta]. \tag{19}$$

Thus, from (18), the 3-D DCT vectors in the F color space lie *approximately* on a parallelepiped lattice $\mathcal{L}_i$ (see Fig. 5(c)). (The DC coefficients lie approximately on a parallelepiped structure that is shifted from the origin.)

6.4 Blind CHEst and nearly orthogonal bases

A JPEG-decompressed image's i^{th}-frequency ($i \neq$ DC) 3-D DCT vectors lie approximately on a lattice $\mathcal{L}_i$ with basis $\mathcal{T}\mathcal{Q}_i$. A key step in blind CHEst is to estimate $\mathcal{T}\mathcal{Q}_i$. However, since a lattice can have multiple bases, we must exploit some additional information about practical affine color transforms to resolve the basis ambiguity.

Practical affine color transforms simply try to find a shifted and approximately rotated reference coordinate system to describe color. Consequently, the linear component $\mathcal{T}$ (and thereby, $\mathcal{T}\mathcal{Q}_i$) of all practical affine color transforms will be "nearly orthogonal". (To be precise, we will assume that $\mathcal{T}$ is *weakly* $\frac{\pi}{3} + \epsilon$-*orthogonal*; see Section 7.3 for the definition.) Therefore, in addition to lattice algorithms, we will also need to understand the properties of nearly orthogonal lattice bases.

7 Lattices Algorithms and Properties of Nearly Orthogonal Bases

We briefly review some celebrated CHEst-relevant lattice problems and some recent results on nearly orthogonal bases.

7.1 Lattice reduction and the Lenstra-Lenstra-Lovasz (LLL) algorithm

A celebrated problem of interest to us is the *lattice reduction* problem, which can be stated as follows: Given a set of vectors b_i's that span a lattice $\mathcal{L}$, find an ordered set of *basis* vectors for $\mathcal{L}$ such that [17]

1. The basis vectors are nearly orthogonal.
2. The shorter basis vectors appear first in the ordering.

Lattice reduction is clearly relevant because we also seek nearly orthogonal lattice bases in CHEst.

A major breakthrough in lattice theory was the discovery of a polynomial time lattice reduction algorithm by Lenstra, Lenstra, and Lovasz [13]; this algorithm is commonly referred to as the *LLL algorithm*. The LLL algorithm can be intuitively understood as an adaptation of Gram-Schmidt orthogonalization [18] that sequentially processes the vectors b_i and maintains a basis spanning the processed vectors. We will invoke LLL in Section 8 to estimate a nearly orthogonal basis that spans the DCT coefficient lattice.

7.2 Closest vector problem (CVP) and Shortest vector problem (SVP)

The *closest vector problem* (CVP) and the *shortest vector problem* (SVP) are two other famous, NP-hard CHEst-relevant lattice problems [14, 15, 19]. They are both closely related to the lattice reduction problem. The CVP aims to find the closest (in the Euclidean sense) lattice point to a given point. For a comprehensive semi-tutorial paper on the CVP and algorithms to solve it, we refer the reader to [14]. The SVP aims to find the shortest non-zero lattice point. The 3-D DCT vectors lie only approximately on a lattice (see Section 6.3). To estimate a basis that approximately span these perturbed DCT vectors, we will invoke CVP solutions in Section 8.2.1.

7.3 Properties of nearly orthogonal lattice basis vectors

Recently, [7, 8] quantified the "orthogonality" of a basis in terms of the angle between its constituent vectors. An ordered set of vectors $\{b_1, b_2, \ldots, b_m\}$ is *weakly θ-orthogonal* if for any $i = 2, \ldots, m$, the angle between b_i and the subspace spanned by $\{b_1, \ldots, b_{i-1}\}$ lies in the range $[\theta, \pi/2]$; that is,

$$\cos^{-1}\left(\frac{|\langle b_i, \sum_{j=1}^{i-1} \alpha_i\, b_i\rangle|}{\|b_i\|_2 \left\|\sum_{j=1}^{i-1} \alpha_i\, b_i\rangle\right\|_2}\right) \geq \theta, \text{ for all } \alpha_j \in \mathbb{R} \text{ with } \sum_j |\alpha_j| > 0. \tag{20}$$

For example, the *ITU.BT-601 YCbCr* to *RGB* transform (see (3)) linear component is weakly θ-orthogonal with $\theta = 0.349\pi$ radians.

Theorem 1 *Any weakly $\frac{\pi}{3} + \epsilon$-orthogonal, $0 < \epsilon \leq \pi/6$, basis contains every shortest non-zero lattice vector.*

Theorem 1's proof, which is provided in [7, 8], follows by induction.

The next theorem addresses the uniqueness of weakly $\frac{\pi}{3} + \epsilon$-orthogonal lattice basis in $\mathbb{R}^3$.

Theorem 2 *Let $\mathcal{B}_1 = \{b_1, b_2, b_3\}$ and $\mathcal{B}_2 = \{\widetilde{b}_1, \widetilde{b}_2, \widetilde{b}_3\}$ be two weakly $\frac{\pi}{3} + \epsilon$-orthogonal bases for a lattice $\mathcal{L}$ in $\mathbb{R}^3$. Let $\mathcal{U}$ be a unimodular matrix such that $\mathcal{B}_1 = \mathcal{B}_2\mathcal{U}$. Then, the $\mathcal{U}$'s elements' absolute values are*

upper-bounded by

$$\kappa(\mathcal{B}_1) := \frac{4}{3}\frac{\max_{j\in\{1,2,3\}} \|b_j\|_2}{\min_{j\in\{1,2,3\}} \|b_j\|_2}. \tag{21}$$

See [8] for the proof and further details. Theorem 2 guarantees that in $\mathbb{R}^3$, a weakly orthogonal basis with nearly equal length vectors is related to every weakly orthogonal basis by a unimodular matrix with small elements. For example, if $\mathcal{B}_1$ is weakly $\frac{\pi}{3}+\epsilon$-orthogonal and its column lengths are within a factor of 1.5 of each other, then the unimodular matrix elements relating $\mathcal{B}_1$ to another weakly $\frac{\pi}{3}+\epsilon$-orthogonal basis $\mathcal{B}_2$ are either 0 or ± 1.

8 Blind Lattice-based CHEst for Color Images

In Section 8.1, we use ideas from Section 7 to develop a blind CHEst approach. Section 8.2 specifies the modifications required to robustify blind CHEst to round-off noise.

8.1 Lattice-based CHEst in the absence of round-off noise

Lattice-based CHEst employs the following steps to solve the blind CHEst problem. Each step's detail is described in the subsections that follow immediately.

1. Estimate weakly $\frac{\pi}{3}$-orthogonal bases $\mathcal{B}_i$'s for AC DCT coefficient lattices $\mathcal{L}_i$'s using LLL.
2. Estimate the color transform's scaled linear component $\mathcal{T}\mathcal{Q}_i$ from the estimated $\mathcal{B}_i$'s.
3. Separate the $\mathcal{T}\mathcal{Q}_i$'s into $\mathcal{T}$ and $\mathcal{Q}_i$'s.
4. Estimate the $\mathcal{Q}_i$ for the DC frequency and color transform's additive component

We will assume that $\mathcal{T}$ is weakly $\frac{\pi}{3}+\epsilon$-orthogonal, $0 \leq \epsilon \leq \frac{\pi}{6}$; see Section 6.4 for this assumption's motivation. We have verified that all the affine color transforms in the literature [1] satisfy this assumption.[4]

8.1.1 Estimating $\frac{\pi}{3}$-orthogonal bases $\mathcal{B}_i$'s using LLL

In the absence of round-offs, any i^{th}-frequency ($i \neq$ DC) 3-D DCT vector $X_{F,i} \in \mathcal{L}_i$ with a $\frac{\pi}{3}$-orthogonal basis $\mathcal{T}\mathcal{Q}_i$. We seek to estimate a weakly $\frac{\pi}{3}$-orthogonal basis $\mathcal{B}_i$ from all the $X_{F,i}$'s. Given the $X_{F,i}$'s $\in \mathcal{L}_i$

[4]Theoretically, our approach can be modified to accommodate cases where $\mathcal{T}$ is not weakly $\frac{\pi}{3}+\epsilon$-orthogonal. However, as the $\mathcal{T}$'s orthogonality decreases, that approach's computational demands would increase and its stability would deteriorate.

as inputs, LLL returns a reduced basis that spans all the $X_{F,i}$'s. (The LLL-reduced basis is non-singular when the input $X_{F,i}$'s span $\mathcal{L}_i$; we will index these frequencies by $i \in \{1, \ldots, p\}$.) The LLL-reduced basis is not guaranteed to be weakly $\frac{\pi}{3}$-orthogonal according to known worst-case bounds on LLL's performance.[5] If the LLL-reduced basis is not $\frac{\pi}{3}$-orthogonal, then we would need to search for a unimodular matrix $\mathcal{U}$ such that $\mathcal{B}_i :=$ LLL-reduced basis times $\mathcal{U}$ is $\frac{\pi}{3}$-orthogonal. However, in our experience, this search has been unnecessary. The LLL output has always been been weakly $\frac{\pi}{3}$-orthogonal for our lattices, which are "well-posed" in the sense that they contain at least one weakly $\frac{\pi}{3}$-orthogonal basis. Our experience conforms with common knowledge that the LLL perform significantly better in practice than what is guaranteed theoretically [16, 17].

8.1.2 Estimating $\mathcal{T}\mathcal{Q}_i$'s

Since $\mathcal{B}_i$ and $\mathcal{T}\mathcal{Q}_i$ are both bases for $\mathcal{L}_i$, we have

$$\mathcal{B}_i = \mathcal{T}\mathcal{Q}_i\mathcal{U}_i. \tag{22}$$

for some unimodular matrix $\mathcal{U}_i$ (not necessarily the identity matrix). Hence estimating the $\mathcal{T}\mathcal{Q}_i$'s from the $\mathcal{B}_i$'s is equivalent to decoding the respective $\mathcal{U}_i$'s.

Thanks to the problem's structure, all the $\mathcal{U}_i$'s satisfy the following constraints:

1. *$\mathcal{B}_i\mathcal{U}_i^{-1}$ is weakly $\frac{\pi}{3}$-orthogonal.*
   This follows from (22).

2. *The columns of $\mathcal{U}_i$ corresponding to $\mathcal{B}_i$'s shortest columns are the standard unit vectors times ± 1.*
   This follows from Theorem 1. Since both $\mathcal{B}_i$ and $\mathcal{T}\mathcal{Q}_i$ are weakly $\frac{\pi}{3}$-orthogonal, they indeed contain the shortest vectors in the $\mathcal{L}_i$.

3. *The product $\mathcal{U}_i\mathcal{B}_i^{-1}\mathcal{B}_j\mathcal{U}_j^{-1}$ is a positive diagonal matrix for any $i, j \in \{1, \ldots, p\}$.*
   This follows from (22). Let $\widetilde{\mathcal{U}}_i$'s, $i \in \{1, \ldots, p\}$, be arbitrary unimodular matrices that satisfy this property for every pair $i, j \in \{1, \ldots, p\}$. It is easy to show [8] that if $\widetilde{\mathcal{U}}_j \neq \mathcal{U}_j$ for some j, then $\widetilde{\mathcal{U}}_i \neq \mathcal{U}_i$ for *every* $i \in \{1, \ldots, p\}$. It follows that to correctly estimate all the $\mathcal{U}_i$'s, we just need to correctly estimate any one $\mathcal{U}_i$. Further, multiple unimodular sequences can satisfy this constraint only when the $\mathcal{Q}_i$'s are chosen carefully.[6]

[5] For 2-D lattices, the LLL is guaranteed to return a weakly $\frac{\pi}{3}$-orthogonal basis.

[6] For example, this is possible if all the $\mathcal{Q}_i$'s are $\mathcal{Q}_1$'s scaled versions and the actual $\mathcal{U}_i$'s are all identity matrices. Then, any

4. *All non-zero elements of $\mathcal{U}_i$ are "small".*

   From Theorem 2, the $\mathcal{U}_i$ elements' absolute values of are $< \kappa(\mathcal{B}_i)$.

Theoretically, the above conditions are not sufficient to uniquely determine the $\mathcal{U}_i$'s. However, even with a small number p of non-singular $\mathcal{B}_i$'s, the above conditions become so restrictive (particularly, constraint 3) that only the actual $\mathcal{U}_i$'s satisfy all them simultaneously in practice. Hence we just need to search for a unimodular matrix sequence that satisfy the above constraints.

The search focuses on estimating the correct unimodular matrix $\mathcal{U}_\ell$ for frequency $\ell := \arg\min_{i\in\{1,\ldots,p\}} \kappa(\mathcal{B}_i)$. Let $\widetilde{\mathcal{U}}_\ell$ denote an arbitrary unimodular matrix that satisfies conditions 1, 2, and 4 for frequency ℓ. From Theorem 2, the frequency ℓ contains the least number of valid $\widetilde{\mathcal{U}}_\ell$'s (along with the correct $\mathcal{U}_\ell$), which makes the search for the correct $\mathcal{U}_\ell$ easier. We sequentially test each valid $\widetilde{\mathcal{U}}_\ell$ and verify if its choice allows us to find valid unimodular matrices $\widetilde{\mathcal{U}}_i$'s for every frequency $i \in \{1,\ldots,p\}$. Note that given $\widetilde{\mathcal{U}}_\ell$, there exists at most one unimodular matrix $\widetilde{\mathcal{U}}_i$ such that $\widetilde{\mathcal{U}}_\ell \mathcal{B}_\ell^{-1} \mathcal{B}_i \widetilde{\mathcal{U}}_i^{-1}$ is positive diagonal; such a $\widetilde{\mathcal{U}}_i$ can be found easily, if it exists. Since the unimodular matrix conditions are extremely restrictive, in practice, we can safely assume that we will be able to find a valid $\widetilde{\mathcal{U}}_i$'s for all frequencies only if $\widetilde{\mathcal{U}}_\ell = \mathcal{U}_\ell$. Further, if $\widetilde{\mathcal{U}}_\ell = \mathcal{U}_\ell$, then $\widetilde{\mathcal{U}}_i = \mathcal{U}_i$. Thus, we can quickly and reliably determine the desired sequence of $\mathcal{U}_i$'s from the $\mathcal{B}_i$'s.

8.1.3 Decomposing $\mathcal{T}\,\mathcal{Q}_i$ into $\mathcal{T}$ and $\mathcal{Q}_i$

Decomposing the $\mathcal{T}\,\mathcal{Q}_i$'s into $\mathcal{T}$ and $\mathcal{Q}_i$'s is equivalent (apart from the sign) to determining the norm of each column of $\mathcal{T}$ because the $\mathcal{Q}_i$'s are diagonal matrices. (The signs of $\mathcal{T}$'s column are chosen such that its largest magnitude entry is positive.) Let $\mathcal{T}(:,k)$ and $(\mathcal{T}\,\mathcal{Q}_i)\,(:,k)$, $k = \{1,2,3\}$, denote $\mathcal{T}$'s and $\mathcal{T}\,\mathcal{Q}_i$'s k^{th} column vectors respectively. Then,

$$\left\|(\mathcal{T}\,\mathcal{Q}_i)\,(:,k)\right\|_2 = q_{i,k}\left\|\mathcal{T}(:,k)\right\|_2. \tag{23}$$

Since the $\mathcal{Q}_i$'s diagonal elements $q_{i,k} \in \mathbb{Z}$, all elements of the set $\{\|(\mathcal{T}\,\mathcal{Q}_i)\,(:,k)\|_2 \,|i \in \{1,\ldots,p\}\}$ lie on the same 1-D lattice. The length of the shortest vector in this lattice is $\|\mathcal{T}(:,k)\|_2$. Hence we set $\mathcal{T}$'s column norm to be the length of the shortest non-zero vector in the 1-D lattice comprising the $\|(\mathcal{T}\,\mathcal{Q}_i)\,(:,k)\|_2$'s.

sequence $\{\widetilde{\mathcal{U}}_i\}$ with $\widetilde{\mathcal{U}}_i = \widetilde{\mathcal{U}}_1$, $i \in \{2,\ldots,p\}$ satisfies constraint 3. Note that to make $\{\widetilde{\mathcal{U}}_i\}$ simultaneously satisfy the other three constraints, $\mathcal{Q}_1$ must be chosen carefully.

8.1.4 Estimating the DC quantization step-sizes and the color transform's additive component

The DC coefficients in the observed space, after being transformed by the estimated $\mathcal{T}$'s inverse, is related to compression color space DC coefficients by (see (3))

$$\mathcal{T}^{-1}X_{F,DC} = \overline{X}_{G,DC} - \text{DCT of additive component}. \tag{24}$$

Since $\overline{X}_{G,DC}$ lies on a rectangular box lattice, $\mathcal{T}^{-1}X_{F,DC}$ lies on a shifted rectangular box lattice. The DC quantization step-size can be obtained by first subtracting an arbitrary reference vector $\mathcal{T}^{-1}X_{F,DC}$ from all the vectors to nullify the shift and then solving the SVP problem along each component.

We estimate the additive component by exploiting two constraints. First, the additive component should be such that the $\overline{X}_{G,DC}$ (see (24)) lie on a lattice. Hence at least some $\overline{X}_{G,DC}$ should be zero. Second, all

$$X_{F,DC} + \mathcal{T}\,\text{DCT of additive component}$$

should lie in the 0–255 range after transformation. Many additive component estimates could satisfy the above two criteria. In such cases, we arbitrarily pick one of the solutions. Note that errors in the additive component estimates do not significantly affect applications such as recompression and enhancement since these applications merely shift the DC coefficients in the compression color space.

8.2 Robustification to round-off noise

We now clarify the modifications required to combat round-off noise.

8.2.1 Estimating $\frac{\pi}{3}$-orthogonal bases $\mathcal{B}_i$'s using a robustified LLL algorithm

We desire to estimate a nearly orthogonal basis $\widehat{\mathcal{B}}_i$ such that all the 3-D DCT vectors $\widetilde{X}_{F,i}$'s lie close to the lattice spanned by $\widehat{\mathcal{B}}_i$ (see (18)). The conventional LLL is unstable when the input vectors $\widetilde{X}_{F,i}$'s contain noise. Hence to estimate $\widehat{\mathcal{B}}_i$, we stabilize LLL in two ways.

First, we exploit the multiplicity of noisy lattice vector realizations to reduce round-off noise propagation. We observe multiple noisy realizations of the lattice vectors because even reasonable-sized images contain many 8×8 pixel blocks. We input the least noisy vectors first to LLL. To determine this input order, we compute the histograms of all the i^{th}-frequency 3-D DCT vectors, and sort the vectors in the descending order of histogram values.

Second, we incorporate a least-squares noise attenuation step that is exploits the round-off errors' statistical distribution. Given a lattice basis estimate $\widetilde{\mathcal{B}}_i$, we assume that any vector lying within a distance ζ (chosen adaptively from the range $[3.5, 5]$) from its closest point on the lattice spanned by $\widetilde{\mathcal{B}}_i$ is just a noisy realization. We use the noisy realization to update the $\widetilde{\mathcal{B}}_i$ and obtain $\widehat{\mathcal{B}}_i$. Let $\mathcal{D}_i$ denote a $3 \times m$ matrix containing the $\widetilde{X}_{F,i}$'s as its columns. Let $\widetilde{\mathcal{I}}_i \in \mathbb{Z}^{3\times m}$ be such that each column of $\widetilde{\mathcal{B}}_i \widetilde{\mathcal{I}}_i$ is the closest point on the lattice spanned by $\widetilde{\mathcal{B}}_i$ to the corresponding column of $\mathcal{D}_i$; that is,

$$\widetilde{\mathcal{B}}_i \widetilde{\mathcal{I}}_i = \text{CVP}_{\widetilde{\mathcal{B}}_i}\left(\mathcal{D}_i\right), \quad \widetilde{\mathcal{I}}_i \in \mathbb{Z}^{3\times m}. \tag{25}$$

Assuming that the distance from $\mathcal{D}_i$ to the closest lattice point is less than ζ, we have using (19)

$$P\Big(\mathcal{D}_i|\widetilde{\mathcal{B}}_i\Big) \propto \exp\left(-\frac{1}{2\sigma^2}\left\|\mathcal{D}_i - \widetilde{\mathcal{B}}_i\widetilde{\mathcal{I}}_i\right\|_{HS}^2\right), \tag{26}$$

where $\|.\|_{HS}^2$ denotes the squared *Hilbert-Schmidt* norm (the square root of the sum of all the matrix elements' squares). We can update the basis estimate as

$$\widehat{\mathcal{B}}_i := \arg\min_{\widetilde{\mathcal{B}}_i} \left\|\mathcal{D}_i - \widetilde{\mathcal{B}}_i\widetilde{\mathcal{I}}_i\right\|_{HS}^2 = \left(\mathcal{D}_i\widetilde{\mathcal{I}}_i^T\right)\left(\widetilde{\mathcal{I}}_i\widetilde{\mathcal{I}}_i^T\right)^{-1}. \tag{27}$$

The above equation assumes that the round-off error norms stay less than ζ and hence ignores the distribution's finite support. The estimation (27) naturally leads to an iterative update where $\widetilde{\mathcal{I}}$ is recomputed using (25) with $\widetilde{\mathcal{B}}_i = \widehat{\mathcal{B}}_i$. As desired, this iteration is guaranteed to converge to a locally optimal $\widehat{\mathcal{B}}_i$ that minimizes the round-off error between the observations $\mathcal{D}_i$ and the closest points on the lattice spanned by $\widehat{\mathcal{B}}_i$.[7]

In summary, we fuse LLL with noise attenuation as follows:

1. Compute the histogram of the AC frequency vectors and sort them in descending order of histogram values.

2. Include the first vector outside the radius ζ as a lattice basis vector. For each frequency i, the $\zeta \in [3.5, 5]$ is set adaptively.) Any vector within the sphere could potentially be a noisy realization of the origin $[0,\ 0,\ 0]^T$, and hence should be ignored.

3. Compute the error vector between the next AC frequency vector and the closest vector on the lattice (obtained by solving a CVP) spanned by the current basis estimate. If the error vector's norm is greater

[7] Convergence follows because both (25) and (27) monotonically reduce $\left\|\mathcal{D}_i - \widetilde{\mathcal{B}}_i\widetilde{\mathcal{I}}_i\right\|_{HS}^2$.

than ζ, then include the currently chosen vector to list of basis vectors, and perform LLL on this set of basis vectors. If the error vector norm is less or equal to ζ, then update the basis vectors using (27).

Combining the update step with LLL successfully curbs the propagation and amplification of the round-off errors during LLL's arithmetic operations.

8.2.2 Estimating the $\mathcal{T}\mathcal{Q}_i$'s

In the absence of round-offs, there exist $\mathcal{U}_i$'s such that $\mathcal{U}_i\mathcal{B}_i^{-1}\mathcal{B}_j\mathcal{U}_j^{-1}$ is exactly diagonal (see Section 8.1.2). However, due to round-offs, $\mathcal{U}_i\widehat{\mathcal{B}}_i^{-1}\widehat{\mathcal{B}}_j\mathcal{U}_j^{-1}$ can only be diagonally dominant. We define the *diagonality* of a matrix as the ℓ_2 norm of the matrix's diagonal elements divided by the ℓ_2 norm of all the matrix elements; the measure is equal to one if and only if the matrix is exactly diagonal. We estimate unimodular matrices $\widehat{\mathcal{U}}_i$'s such that sum of the $\widehat{\mathcal{U}}_i\widehat{\mathcal{B}}_i^{-1}\widehat{\mathcal{B}}_j\widehat{\mathcal{U}}_j^{-1}$'s diagonality measures is maximized. We set $\widehat{\mathcal{T}}\,\widehat{\mathcal{Q}}_i := \widehat{\mathcal{B}}_i\widehat{\mathcal{U}}_i^{-1}$, with $\widehat{\mathcal{T}}$ and $\widehat{\mathcal{Q}}_i$ denoting $\mathcal{T}$'s and $\mathcal{Q}_i$'s estimates respectively.

8.2.3 Estimating $\mathcal{T}$ and $\mathcal{Q}_i$'s

The estimated $\widehat{\mathcal{T}}\,\widehat{\mathcal{Q}}_i$'s column norms conform only approximately to a 1-D lattice spanned by the corresponding column norm of the true $\mathcal{T}$. Similar to the quantization step-size estimation described in Section 4, we estimate $\widehat{\mathcal{T}}$'s column norms by solving a penalized least-squares cost function

$$\left\|\widehat{\mathcal{T}}(:,k)\right\|_2 = \arg\min_{\varrho} \sum_i \left(\left(\left\| \left(\widehat{\mathcal{T}}\,\widehat{\mathcal{Q}}_i\right)(:,k) \right\|_2 - \varrho\,\mathrm{round}\left(\frac{\left\| \left(\widehat{\mathcal{T}}\,\widehat{\mathcal{Q}}_i\right)(:,k) \right\|_2}{\varrho} \right) \right)^2 + \beta\,\mathrm{round}\left(\frac{\left\| \left(\widehat{\mathcal{T}}\,\widehat{\mathcal{Q}}_i\right)(:,k) \right\|_2}{\varrho} \right) \right). \tag{28}$$

The first term ensures that $\widehat{\mathcal{T}}\,\widehat{\mathcal{Q}}_i$'s column norms conform to a 1-D lattice spanned by ϱ, and the second term ensures that $\widehat{\mathcal{T}}$'s column norm is large. The β controls the tradeoff between the two terms. In practice, we set $\beta := \frac{0.2}{\mathrm{mean}\left(\|(\mathcal{T}\mathcal{Q}_i)(:,k)\|_2\right)}$. We can then estimate the quantization step-sizes for all the AC frequencies as

$$\widehat{q}_{i,k} = \mathrm{round}\left(\frac{\left\| \left(\widehat{\mathcal{T}}\,\widehat{\mathcal{Q}}_i\right)(:,k) \right\|_2}{\left\|\widehat{\mathcal{T}}(:,k)\right\|_2} \right).$$

8.2.4 Estimating the additive component and the DC quantization step-sizes

Each component of any arbitrary$\widehat{\mathcal{T}}^{-1}\widetilde{X}_{F,DC}$ is approximately equal to an integer multiple of the respective DC component's quantization step-size plus a constant shift. Hence, the histogram of each component of the $\widehat{\mathcal{T}}^{-1}\widetilde{X}_{F,DC}$ collection looks like a shifted version of Fig. 3. We note that the histogram's Discrete Fourier Transform (DFT) magnitude is immune to the unknown constant shift. Further, the DFT's peak frequency captures the histogram's "periodicity", which is determined by the quantization step-size. Hence, we estimate the quantization step-size as the inverse of the non-zero frequency at which the histogram's DFT magnitude peaks. Subsequently, we estimate the affine transform's additive component as described in Section 8.1.

8.3 Lattice-based CHEst results

We demonstrate the performance on lattice-based CHEst on the 512×512 *Lena* color image [12] that we JPEG-compressed in the *ITU.BT-601 YCbCr* space (see (3)). Lattice-based CHEst performed equally well on a wide variety of other experiments comprising different images and compression color spaces. The luminance plane Y's DCT coefficients were quantized using table 1 from Fig. 2 and the chrominance planes Cb's and Cr's DCT coefficients were quantized using table 2 from Fig. 2. The *Cb* and *Cr* planes were not subsampled during compression. The image was then decompressed and then transformed to the *RGB* space. The algorithm operated in this *RGB* space and tried to estimate the affine transformation from *ITU.BT-601 YCbCr* to the current *RGB* space (see (3)).

Lattice-based CHEst estimated that the affine transform from the compression space *ITU.BT-601 YCbCr* to the observation space *RGB* was

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} = \begin{bmatrix} 1.00 & 0.00 & 1.41 \\ 1.00 & -0.35 & -0.71 \\ 1.00 & 1.78 & 0.00 \end{bmatrix} \left(\begin{bmatrix} Y \\ Cb \\ Cr \end{bmatrix} - \begin{bmatrix} 3 \\ 88 \\ 138 \end{bmatrix} \right). \tag{29}$$

Figure 6 illustrates the algorithm's quantization table estimates. An $\times$ indicates that the quantization step-size estimation was not possible because all DCT coefficients were quantized to zero. The estimated CH conforms well with the true compression settings; compare (3) to (29) and Fig. 2 to Fig. 6.

We now outline the results obtained by the algorithm's various intermediate steps. In the first step (see Step 1 in Section 8.1), robustified LLL (details in Section 8.2.1) estimated the lattice bases for the AC

frequencies $[1, 2]$ and $[1, 3]$ as

$$\widehat{\mathcal{B}}_{[1,2]} = \begin{bmatrix} -7.00 & 15.50 & -6.96 \\ -7.01 & -7.81 & -10.72 \\ -7.00 & 0.02 & 12.66 \end{bmatrix} \quad \text{and} \quad \widehat{\mathcal{B}}_{[1,3]} = \begin{bmatrix} -6.01 & 13.64 & -5.95 \\ -5.99 & -15.91 & -10.90 \\ -6.00 & -5.91 & 18.94 \end{bmatrix}. \tag{30}$$

Clearly, the respective first columns of $\widehat{\mathcal{B}}_{[1,2]}$ and $\widehat{\mathcal{B}}_{[1,3]}$, which are the shortest columns, are indeed aligned with one of the columns of the *ITU.BT-601 YCbCr* to *RGB* transformation's linear component $\mathcal{T}$. However, $\widehat{\mathcal{B}}_{[1,2]}$'s third column and $\widehat{\mathcal{B}}_{[1,2]}$'s second and third column are not scaled versions of any of $\mathcal{T}$'s columns due to the addition of the first column.

In the second step (see Step 2 in Section 8.1), we deduced the unimodular matrices to be

$$\widehat{\mathcal{U}}_{[1,2]} = \begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad \text{and} \quad \widehat{\mathcal{U}}_{[1,3]} = \begin{bmatrix} 1 & 1 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}. \tag{31}$$

Hence, the estimate $\widehat{\mathcal{T}}\,\widehat{\mathcal{Q}}_i = \widehat{\mathcal{B}}_i\widehat{\mathcal{U}}_i^{-1}$ for the AC frequencies $[1, 2]$ and $[1, 3]$ is

$$\widehat{\mathcal{B}}_{[1,2]}\widehat{\mathcal{U}}_{[1,2]}^{-1} = \begin{bmatrix} -7.00 & 15.50 & 0.04 \\ -7.01 & -7.81 & -3.70 \\ -7.00 & 0.02 & 19.66 \end{bmatrix} \quad \text{and} \quad \widehat{\mathcal{B}}_{[1,3]}\widehat{\mathcal{U}}_{[1,3]}^{-1} = \begin{bmatrix} -6.01 & 19.65 & 0.06 \\ -5.99 & -9.92 & -4.91 \\ -6.00 & 0.09 & 24.94 \end{bmatrix}. \tag{32}$$

Similarly, we computed the $\widehat{\mathcal{T}}\,\widehat{\mathcal{Q}}_i$ for all the AC frequencies.

In the third step (see Step 3 in Section 8.1), we estimated $\mathcal{T}$'s column norms using (28). This yielded all the AC frequency quantization step-sizes $\widehat{\mathcal{Q}}_i$ illustrated in Fig. 6 and the $\mathcal{T}$ (the 3×3 matrix) in (29).

In the fourth and final step (see Step 4 in Section 8.1), we used the DC frequency coefficients to estimate the additive component as the 3×1 matrix in (29), and the DC quantization step-sizes as shown in Fig. 6.

9 JPEG Recompression: An Example Application of CHEst

When a given TIFF or BMP image's file-size needs to be reduced, the conventional approach is to naively employ JPEG with an arbitrary choice of compression color space, subsampling factor, and quantization table. For naive JPEG recompression, reasonable choices for the color transformations include *RGB* to *ITU.BT-601 YCbCr*, *Computer* RGB to *ITU.BT-601 YCbCr*, *RGB* to *Kodak PhotoYCC*, and *sRGB* to 8-bit

10	7	6	10	14	24	31	×
7	7	8	11	16	35	36	33
8	8	10	14	24	34	×	×
8	10	13	17	31	×	×	×
11	13	22	34	41	×	×	×
14	21	×	×	×	×	×	×
×	×	×	×	×	×	×	×
×	×	×	×	×	×	×	×

Y plane

10	11	14	28	×	×	×	×
11	13	16	×	×	×	×	×
14	16	×	×	×	×	×	×
×	×	×	×	×	×	×	×
×	×	×	×	×	×	×	×
×	×	×	×	×	×	×	×
×	×	×	×	×	×	×	×
×	×	×	×	×	×	×	×

Cb plane

10	11	14	28	×	×	×	×
11	13	16	×	×	×	×	×
14	16	×	×	×	×	×	×
×	×	×	×	×	×	×	×
×	×	×	×	×	×	×	×
×	×	×	×	×	×	×	×
×	×	×	×	×	×	×	×
×	×	×	×	×	×	×	×

Cr plane

Figure 6: *Quantization table estimates using lattice-based CHEst.*

CIELab. Some common subsampling factors are $2 \times 2, 1 \times 1, 1 \times 1$ and $1 \times 1, 1 \times 1, 1 \times 1$. The quantization tables are often set by adjusting a so-called Quality Factor (QF), a reference number between 1 to 100 used by the IJG JPEG implementation [9]; QF=100 sets all the quantizer step-sizes to unity and thus yields the best quality JPEG can possibly achieve. Any combination of the above choices would yield a JPEG image file with a certain file-size. Smaller file-sizes are typically accompanied by increased distortions in the recompressed image. In this section, we demonstrate that using CHEst to recompress a JPEG-decompressed color image offers significant benefits over a naive recompression approach.

9.1 JPEG recompression using dictionary-based CHEst

To demonstrate the dictionary-based CHEst's benefits in JPEG recompression, consider the same test image described in Section 5.3—*Lena* color image previously JPEG-compressed in the *8-bit CIELab* color space with $2 \times 2, 1 \times 1, 1 \times 1$ subsampling using quantization table 1 and 2 from Fig. 2. As described in Section 5.3, dictionary-based CHEst accurately estimates the test image's CH.

To perform recompression using the dictionary-based CHEst information, we first transformed the observed image into the estimated compression color space using the *sRGB* to *8-bit CIELab* color transformation. Then, we deconvolved the effect of the smoothing employed during previous decompression on

the *a* and *b* color planes. After performing $2 \times 2, 1 \times 1, 1 \times 1$ subsampling, using the IJG JPEG implementation [9], we JPEG-compressed the 8-bit CIELab color planes with the estimated quantization tables in Fig. 4 (setting the $\times$ entries to 100). Our recompression yielded a JPEG image with file-size 32.31 kilobytes (KB) with an SNR of 22.58 dB. The SNR is computed in dB with respect to the original *Lena* image in the perceptually-uniform CIELab color space. We also visually inspected the images to confirm that the SNR values were consistent with the image's visual quality.

For comparison, we also recompressed the image using a variety of naively chosen settings. We JPEG-compressed the test BMP image using the *RGB* to *ITU.BT-601 YCbCr*, *Computer RGB* to *ITU.BT-601 YCbCr*, *RGB* to *Kodak PhotoYCC*, and *sRGB* to 8-bit CIELab color transforms using $2 \times 2, 1 \times 1, 1 \times 1$ and also using $1 \times 1, 1 \times 1, 1 \times 1$ subsampling. For each chosen color transform and subsampling, we varied the quantization tables using the QF value and noted the resulting JPEG image's file-size (in KB) and the incurred distortion in SNR (in dB in the CIELab space).

Figures 7(a) and (b) summarize the recompression results. In both plots, the "$\diamondsuit$" symbol marks the file-size SNR pair (32.31 KB, 22.58 dB) associated with the image recompressed using dictionary-based CHEst results. Each curve in Fig. 7(a) illustrates the achieved file-size versus SNR tradeoff for naive recompression in the indicated color space with $2 \times 2, 1 \times 1, 1 \times 1$ subsampling. The curves in Fig. 7(b) illustrate the tradeoff when $1 \times 1, 1 \times 1, 1 \times 1$ subsampling is employed. The naive recompression curves demonstrate a "knee-point" trend—the SNR remains flat for a broad file-size range, but decreases rapidly for small file-size changes thereafter. (Arguments similar to those in [4] could be used to explain the non-monotonicity of the naive recompression curves.) Both the plots confirm that exploiting the dictionary-based CHEst enables us to strike a desirable file-size versus distortion tradeoff—we attain the nearly minimum file-size without introducing significant additional distortion.

9.2 JPEG recompression using lattice-based CHEst

We demonstrate the lattice-based CHEst's benefits in JPEG recompression using the test image described in Section 8.3—*Lena* color image previously JPEG-compressed in the *ITU.BT-601 YCbCr* color space with $1 \times 1, 1 \times 1, 1 \times 1$ subsampling using quantization tables 1 and 2 from Fig. 2. As described in Section 8.3, lattice-based CHEst accurately estimates the test image's CH.

To perform recompression using the lattice-based CHEst results, we transformed the observed image to the estimated compression space space using the inverse of estimated *ITU.BT-601 YCbCr* to *RGB* transformation in (29). We JPEG-compressed the three planes using the estimated quantization tables in Fig. 2

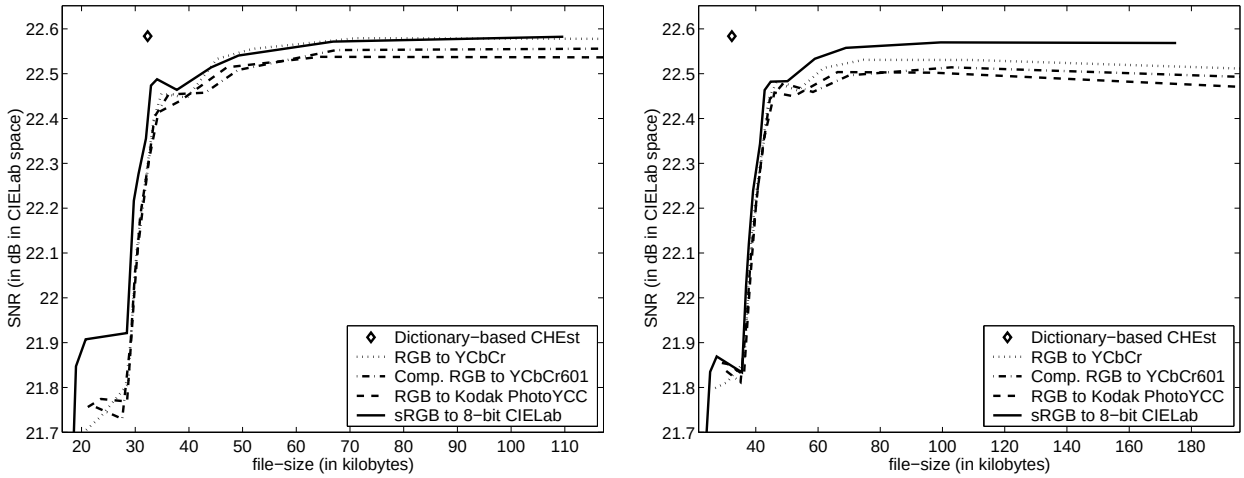

(a) $2 \times 2, 1 \times 1, 1 \times 1$ subsampling (b) $1 \times 1, 1 \times 1, 1 \times 1$ subsampling

Figure 7: *JPEG recompression results for dictionary-based CHEst. The "♢" marks the file-size SNR pair (*32.31 *KB,* 22.58 *dB) obtained using dictionary-based CHEst results for JPEG recompression. Each curve in Fig. 7(a) illustrates the achieved file-size versus SNR tradeoff for naive recompression in the indicated color space with* $2 \times 2, 1 \times 1, 1 \times 1$ *subsampling. Plot (b) illustrates the tradeoff when* $1 \times 1, 1 \times 1, 1 \times 1$ *subsampling is employed.*

(setting the $\times$ entries to 100) to obtain an image with file-size=44.81 KB and SNR=24.03 dB.

Figure 8(a) and (b) compares the file-size SNR pair for lattice-based CHEst recompression with file-size versus SNR curves for naive JPEG recompression in different color spaces at different QFs for $2 \times 2, 1 \times 1, 1 \times 1$ and $1 \times 1, 1 \times 1, 1 \times 1$ subsampling. Figure 8 verifies that lattice-based CHEst results also enables us to strike a desirable file-size versus distortion tradeoff during JPEG recompression.

10 Conclusions

This paper addressed the JPEG CHEst problem for color images and its potential applications. JPEG compression leaves its signature on an image by quantizing the image's DCT coefficients and forcing them to closely conform to near-periodic structures. The paper described two new approaches that exploit these structures to solve the CHEst problem.

First, we formulated a statistical framework to characterize and exploit these JPEG-induced near-periodic structures for gray-scale and color images. Essentially, the statistical approach chooses from a dictionary the best CH model that explains the regular structure of the observed image's DCT coefficients.

Second, for cases when JPEG employs affine color transforms and no subsampling, we devised a blind CHEst scheme that does not rely on a finite dictionary. In this case, the JPEG-decompressed image's DCT

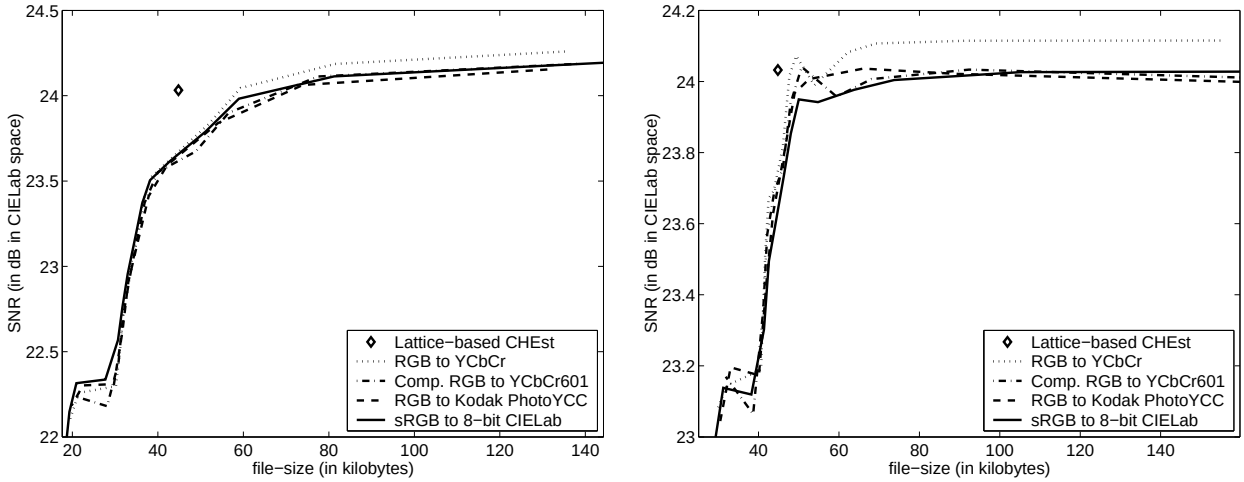

(a) $2 \times 2, 1 \times 1, 1 \times 1$ subsampling (b) $1 \times 1, 1 \times 1, 1 \times 1$ subsampling

Figure 8: *Recompression results for lattice-based CHEst. Recompression using lattice-based CHEst information yields a JPEG image whose file-size is* 44.81 *KB and SNR is* 24.03 *dB; a "♢" marks this file-size SNR pair. Similar to Fig. 7, the curves (a) and (b) illustrate the achieved file-size versus SNR tradeoff for naive recompression in the indicated color space with* $2 \times 2, 1 \times 1, 1 \times 1$ *and* $1 \times 1, 1 \times 1, 1 \times 1$ *subsampling respectively.*

coefficients conform to 3-D lattice structures. The JPEG CH information is encoded in the nearly orthogonal bases that span the DCT lattices. By exploiting recent insights on nearly orthogonal lattice bases and existing lattice algorithms, we provided a novel blind lattice-based solution to the CHEst problem.

JPEG CHEst offers significant benefits during the recompression of JPEG-decompressed color images compared to a naive approach. We demonstrated that exploiting the estimated CH during JPEG recompression introduces minimal distortion (large signal-to-noise-ratio) and simultaneously achieves a small file-size.

JPEG CHEst could also help us uncover proprietary compression settings used by imaging devices. It could contribute to applications such as covert message passing and image authentication. In summary, we envision that JPEG CHEst would enable a variety of intriguing applications.

Acknowledgments

Thanks to Matt Gaubatz for helping us with the color transform implementations.

References

[1] C. Poynton, *A Technical Introduction to Digital Video*. New York: Wiley, 1996.

[2] G. Sharma and H. Trussell, "Digital color imaging," *IEEE Trans. Image Processing*, vol. 6, pp. 901–932, July 1997.

[3] W. Pennebaker and J. Mitchell, *JPEG, Still Image Data Compression Standard*. Van Nostrand Reinhold, 1993.

[4] H. H. Bauschke, C. H. Hamilton, M. S. Macklem, J. S. McMichael, and N. R. Swart, "Recompression of JPEG images by requantization," *IEEE Trans. Image Processing*, vol. 12, pp. 843–849, Jul. 2003.

[5] Z. Fan and R. Eschbach, "JPEG decompression with reduced artifacts," in *Proc. IS&T/SPIE Symp. Electronic Imaging: Image and Video Compression*, (San Jose, CA), Feb. 1994.

[6] 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.

[7] R. Neelamani, *Inverse Problems in Image Processing*. Ph.D. dissertation, ECE Dept., Rice University, 2003. `www.dsp.rice.edu/~neelsh/publications`.

[8] R. Neelamani, S. Dash, and R. Baraniuk, "On nearly orthogonal lattice bases," *IBM Technical Report*, 2004. `www.dsp.rice.edu/~neelsh/publications`.

[9] *Independent JPEG Group Library*. `www.ijg.org`.

[10] R. J. Clarke, *Transform Coding of Images*. London, England: Academic Press, 1985.

[11] A. N. Tikhonov and V. Y. Arsenin, *Solutions of Ill-Posed Problems*. Washington D.C.: V. H. Winston & Sons, 1977.

[12] "The USC-SIPI image database." `sipi.usc.edu/services/database/Database.html`.

[13] A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász, "Factoring polynomials with rational coefficients," *Mathematics Annalen*, vol. 261, pp. 515–534, 1982.

[14] E. Agrell, T. Eriksson, A. Vardy, and K. Zeger, "Closest point search in lattices," *IEEE Trans. Inform. Theory*, vol. 48, pp. 2201–2214, Aug. 2002.

[15] R. Kannan, "Algorithmic geometry of numbers," *Annual Review of Computer Science*, vol. 2, pp. 231–267, 1987.

[16] P. Nguyen and J. Stern, "Lattice reduction in cryptology: An update," in *Lecture notes in Comp. Sci.*, vol. 1838, pp. 85–112, Springer Verlag, 2000.

[17] A. Joux and J. Stern, "Lattice reduction: A toolbox for the cryptanalyst," *Journal of Cryptology*, vol. 11, no. 3, pp. 161–185, 1998.

[18] G. H. Golub and C. F. V. Loan, *Matrix Computations*. Baltimore: Johns Hopkins University Press, 1989.

[19] M. Ajtai, "The shortest vector problem in L_2 is NP-hard for randomized reductions," in *Thirtieth Annual ACM Symposium on Theory of Computing, ACM Press*, pp. 10–19, 1998.