HIGH-DYNAMIC RANGE COMPRESSION USING A FAST MULTISCALE OPTIMIZATION

*Matthieu Maitre**

University of Illinois at Urbana-Champaign
Coordinated Science Laboratory
Urbana, IL

Yunqiang Chen, Fang Tong

Siemens Corporate Research
Intelligent Vision and Reasoning
Princeton, NJ

ABSTRACT

High-dynamic-range medical images have intensity values which cannot be visualized on current low-dynamic-range displays. In this paper, we introduce a fast range compression method which avoids common artifacts such as loss of contrast, haloes and gradient inversions. The proposed method first compresses the intensity range using a global transfer function. It then extracts and enhances weak structures using a multiscale decomposition and artifact correction. We show that artifact correction can be formulated as a linear-programming problem, for which we propose an efficient approximate solution. Experiments on real data demonstrate the effectiveness and speed of the proposed algorithm.

***Index Terms*—** High-dynamic-range images, range compression, artifact correction, Laplacian pyramid, linear programming

1. INTRODUCTION

Medical imaging devices are able to generate data whose High-Dynamic Range (HDR) far exceeds the display capabilities of current low-dynamic-range monitors. It is therefore necessary to compress their range before visualizing them. The challenge here is to design techniques which reduce the intensity range but preserve the image structural content, so that as much information as possible be available for diagnostics.

Range-compression methods fall into two categories [1]: tone-reproduction curves and tone-reproduction operators. Tone-reproduction curves apply the same global transfer function independently to all pixels. They are usually faster than tone-reproduction operators. However, they are also less flexible, which limits the quality of the compressed images. Such methods include simple transfer functions like logarithm, power, or linear functions. Improved results are obtained by defining the transfer function based on intensity histograms [2] or by relying on more complex transfer functions [3]. Tone-reproduction operators take into account the pixels along with their local neighborhoods. Such methods include direct intensity processing [4], two-band decompositions [5], multi-band decompositions [6], and gradient-domain processing [7].

Range-compression methods suffer from severe artifacts, such as loss of contrast [7], haloes around edges [6], or gradient reversals in slowly varying regions [8]. All of these artifacts make diagnostics more difficult: reduced contrasts remove weak but meaningful structures, while haloes and gradient reversals actually add distracting structures to the images. Moreover, the most successful methods tend to be computationally demanding [7], which hinders their implementation in medical devices.

*The author performed the work while at Siemens Corporate Research.

In this paper, we present a novel algorithm which aims at compressing HDR images while preserving their structural content. The algorithm first compresses the intensity range using a global transfer function. Since it reduces the image contrast, an enhancement procedure follows which amplifies the weak structures to keep them visible. Here, we rely on a Laplacian pyramid [9] to extract and enhance these structures. We present a novel theoretical analysis of the enhancement artifacts, which leads to a linear-programming problem [10]. Since the size of the images makes linear-programming solvers too complex to be practical, we propose an efficient method which provides an approximate solution. The proposed algorithm is able to quickly process the large images found in typical medical scenarii, which makes it suitable for medical devices. Experiments on real data show the effectiveness and speed of this algorithm.

The remainder of the paper proceeds as follows. First, Section 2 provides a theoretical analysis of the enhancement problem and its artifacts. It leads to a linear-programming formulation, for which Section 3 presents an approximate but efficient solution. Finally, Section 4 describes our experimental results.

2. THEORETICAL ANALYSIS

We now turn to the theoretical study of HDR compression and its artifacts. The goal of HDR compression is to transform an input image s into an output image $\hat{s}$ with a reduced intensity range but a similar structural content.

In the following, the grayscale input image is assumed to be a 2D signal made of non-negative integer values s_{ij}, sampled at the integer locations (i, j) on a rectangular lattice. For readability, we shall omit the subscripts when obvious. The output signal $\hat{s}$ is defined in a similar fashion.

An intermediate compressed signal $\tilde{s}$ is first generated by applying a global transfer function to the image, that is $\tilde{s} = f(s)$. The compressed signal is then decomposed into a Gaussian pyramid. Let us denote by $\uparrow 2$ and $\downarrow 2$ respectively the upsampling and downsampling operators, along with their associated low-pass filters. The bands of the Gaussian pyramid are related by $\tilde{s}^{(l+1)} = \downarrow 2(\tilde{s}^{(l)})$, where $0 \leq l < L$ is the pyramid level. At the finest level $\tilde{s}^{(0)} = \tilde{s}$.

These bands are upsampled to compute the high-pass bands of a Laplacian pyramid. Denoting $\tilde{c}^{(l)}$ the coarse signals obtained by upsampling, that is $\tilde{c}^{(l)} = \uparrow 2(\tilde{s}^{(l+1)})$, the high-pass bands are given by the analysis equation

$$\tilde{d}^{(l)} = \tilde{s}^{(l)} - \tilde{c}^{(l)}. \tag{1}$$

Enhancing the weak structures amounts to modulating the high-pass bands by gain functions $g^{(l)}(.)$, that is

$$\hat{d}^{(l)} = g^{(l)}\Big(\tilde{d}^{(l)}\Big)\ \tilde{d}^{(l)}, \tag{2}$$

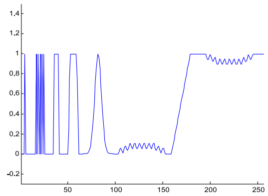

(a) Original signal

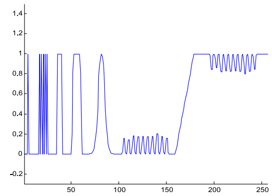

(b) Signal enhanced with artifact correction

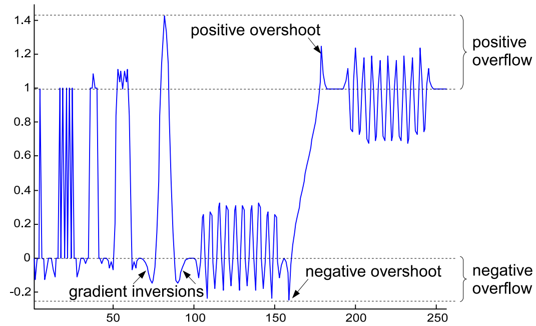

(c) Signal enhanced without artifact correction

Fig. 1: The proposed method corrects the artifacts introduced by weak-structure enhancement, such as gradient inversions, overshoots, and overflows.

and correcting artifacts. The gain functions $g^{(l)}(.)$ can be any functions whose values are greater or equal to one.

The enhanced signals then follow from the synthesis equation

$$\hat{s}^{(l)} = \hat{c}^{(l)} + \hat{d}^{(l)} \tag{3}$$

where $\hat{c}^{(l)}$ denotes the coarse enhanced signal obtained by upsampling the enhanced signal $\hat{s}^{(l+1)}$ at the previous levels, that is $\hat{c}^{(l)} = \uparrow 2(\hat{s}^{(l+1)})$. This top-down process is initialized by $\hat{s}^{(L-1)} = \tilde{s}^{(L-1)}$ at the coarsest level.

The weak-structure enhancement and the artifact correction are computed jointly by solving a series of independent optimization problems, one at each level. For clarity, in the following we drop the exponent (l).

The goal is to obtain an enhancement signal $\hat{d}$ which is:

- as close as possible to the one given by the gain equation (2),
- does not exceed it,
- has the same signs as $\tilde{d}$,
- does not create artifacts such as overflows, overshoots, and gradient inversions in the enhanced signal $\hat{s}$ (see Figure 1).

The first three conditions can be expressed by the optimization problem

$$\max_{\hat{d}} \sum_{i,j} |\hat{d}_{ij}| \qquad \text{such that } 0 \leq \chi(\tilde{d}_{ij})\hat{d}_{ij} \leq g(\tilde{d}_{ij})|\tilde{d}_{ij}|,\ \forall(i,j), \tag{4}$$

where $\chi(.)$ denotes a sign operator which takes the value 1 when its input is non-negative, and -1 otherwise.

The last condition is enforced by introducing two additional sets of constraints in the optimization (4). The first set of constraints prevents positive and negative overflows. It is expressed as

$$0 \leq \hat{s} \leq s_{max}, \tag{5}$$

where s_{max} is the maximum output value, e.g. 255 in the case of an 8-bit display. From (3) it follows that

$$-\hat{c} \leq \hat{d} \leq s_{max} - \hat{c}. \tag{6}$$

The second set of constraints tests the partial derivatives of the signal to prevent gradient inversions and overshoots, both positive and negative. Let us denote the partial derivatives along the axes i and j by respectively $\partial/\partial i$ and $\partial/\partial j$, which are approximated by the finite difference kernel $[-1\ 1]$. For the time being, we only consider the partial derivatives along the axis i.

Overshoots are reduced by limiting the increase of the partial-derivative magnitude,

$$\left|\frac{\partial \hat{s}}{\partial i}\right| \leq \beta_{max} \left|\frac{\partial \tilde{s}}{\partial i}\right| \tag{7}$$

where β_{max} is a constant factor greater than one.

Gradient inversions are prevented by enforcing that the partial derivatives of $\tilde{s}$ and $\hat{s}$ have the same sign. Here, we actually rely on a stronger version of this constraint, which also prevents the enhancement signal from completely flattening our the signal,

$$\chi\left(\frac{\partial \tilde{s}}{\partial i}\right)\frac{\partial \hat{s}}{\partial i} \geq \beta_{min} \left|\frac{\partial \tilde{s}}{\partial i}\right| \tag{8}$$

where β_{min} is a constant factor smaller than one.

Since (8) enforces that $\tilde{s}$ and $\hat{s}$ have partial derivatives with the same sign, we have

$$\left|\frac{\partial \hat{s}}{\partial i}\right| = \chi\left(\frac{\partial \hat{s}}{\partial i}\right)\frac{\partial \hat{s}}{\partial i} = \chi\left(\frac{\partial \tilde{s}}{\partial i}\right)\frac{\partial \hat{s}}{\partial i}. \tag{9}$$

Therefore, (7) and (8) can be merged into a unique set of constraints. From (3) it follows that

$$\beta_{min}\left|\frac{\partial \tilde{s}}{\partial i}\right| - \varepsilon \leq \chi\left(\frac{\partial \tilde{s}}{\partial i}\right)\left(\frac{\partial \hat{d}}{\partial i} + \frac{\partial \hat{c}}{\partial i}\right) \leq \beta_{max}\left|\frac{\partial \tilde{s}}{\partial i}\right| + \varepsilon. \tag{10}$$

where the small constant ε has been added to cope with noisy signals. The same constraint holds for the partial derivatives along the j axis.

Putting equations (4), (6) and (10) together, we obtain the following optimization problem,

$$\begin{array}{c}
\max_{\hat{d}} \sum_{i,j} |\hat{d}_{ij}| \\
\text{such that } \forall(i,j), \\
-\hat{c}_{ij} \leq \hat{d}_{ij} \leq s_{max} - \hat{c}_{ij}, \\
0 \leq \chi(\tilde{d}_{ij})\hat{d}_{ij} \leq g(\tilde{d}_{ij})|\tilde{d}_{ij}|, \\
\beta_{min}\left|\frac{\partial \tilde{s}_{ij}}{\partial i}\right| - \varepsilon \leq \chi\left(\frac{\partial \tilde{s}_{ij}}{\partial i}\right)\left(\frac{\partial \hat{d}_{ij}}{\partial i} + \frac{\partial \hat{c}_{ij}}{\partial i}\right) \leq \beta_{max}\left|\frac{\partial \tilde{s}_{ij}}{\partial i}\right| + \varepsilon, \\
\beta_{min}\left|\frac{\partial \tilde{s}_{ij}}{\partial j}\right| - \varepsilon \leq \chi\left(\frac{\partial \tilde{s}_{ij}}{\partial j}\right)\left(\frac{\partial \hat{d}_{ij}}{\partial j} + \frac{\partial \hat{c}_{ij}}{\partial j}\right) \leq \beta_{max}\left|\frac{\partial \tilde{s}_{ij}}{\partial j}\right| + \varepsilon.
\end{array} \tag{11}$$

As is, this optimization is non-linear and non-derivable due to the sum of absolute values in the objective function. However, it can be transformed into a linear-programming problem [10] by splitting the negative and positive parts of the enhancement signal. Denoting them respectively $\hat{d}_-$ and $\hat{d}_+$, the maximization term in (11) is equivalent to

$$\begin{array}{c}
\max_{\hat{d}_+,\hat{d}_-} \sum_{i,j}\left(\hat{d}_{ij+} + \hat{d}_{ij-}\right) \\
\text{such that } \forall(i,j), \\
\hat{d}_{ij} = \hat{d}_{ij+} - \hat{d}_{ij-}, \\
0 \leq \hat{d}_{ij-} \leq s_{max}, 0 \leq \hat{d}_{ij+} \leq s_{max},
\end{array} \tag{12}$$

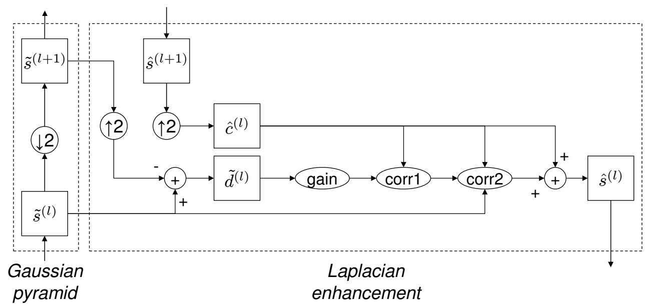

Fig. 2: One level of weak-structure enhancement. A bottom-up process builds the Gaussian pyramid. A top-down process then computes the high-pass bands of a Laplacian pyramid and enhances them using a gain map and two correction passes ('corr1' and 'corr2').

which is linear.

Linear programming problems are particularly interesting. First, there are no poor local optima since all local optima are also global optima [10]. Second, these problems have been extensively studied and several methods exist to solve them [10]. Here, however, the size of the images, and therefore the number of variables and constraints in the optimization, precludes these methods. Instead, we propose an efficient algorithm which provides an approximate solution.

3. IMPLEMENTATION

As mentioned in the previous section, the proposed algorithm first reduces the image range using a global transfer function and then enhances the weak structures by finding an approximate solution to the optimization problem (11). Figure 2 gives an overview of the enhancement process.

The separable binomial filter $[1\ 4\ 6\ 4\ 1]/16$ is used as low-pass filter in the upsampling and downsampling operators of the Laplacian pyramid. The binomial filter does not suffer from ringing artifacts, which means that upsampling and downsampling do not create artifacts. The binomial filter is implemented by multiplication-less lifting [11]. The global transfer function is chosen to be logarithmic and applied using a look-up table.

3.1. Weak-Structure Amplification

The enhancement begins by setting the enhanced signal $\hat{d}$ to the values given by the gain equation (2). This implements the objective function and the second constraint of the optimization problem (11).

Here we rely on the following gain function

$$g(\tilde{d}_{ij}) = \min_{(k,l)\in\mathcal{N}_{3\times 3}(i,j)} \left(1 + \alpha e^{-\frac{\tilde{d}_{kl}^2}{\sigma^2}}\right) \tag{13}$$

where $\mathcal{N}_{3\times3}(i,j)$ is the 3×3 block of pixels centered at (i,j), α is a parameter controlling the enhancement, and σ^2 is a parameter controlling the bias toward weak structures. Both parameters can take different values at each level to enhance specific structures.

3.2. First Correction

The first correction pass reduces the enhancement signal $\hat{d}$ to avoid overflow artifacts. This implements the first constraint of the optimization problem (11). The values of the enhancement signal $\hat{d}$ are

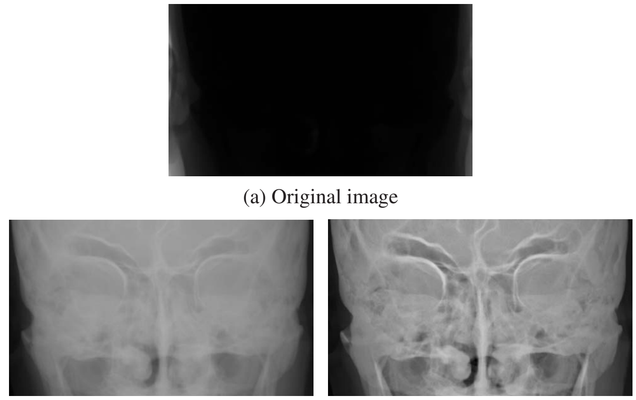

(a) Original image

(b) Global transfer function

(c) Proposed algorithm

Fig. 3: A crop from the skull image (a) processed without enhancement (b) and with enhancement (c). The proposed algorithm is able to compress the range while preserving the image structural content.

updated using the equation

$$\hat{d}_{ij} \leftarrow \begin{cases} \min\left(\hat{d}_{ij}, -\hat{c}_{ij} + s_{max}\right) & \text{if } d_{ij} \geq 0, \\ \max\left(\hat{d}_{ij}, -\hat{c}_{ij}\right) & \text{otherwise.} \end{cases} \tag{14}$$

3.3. Second Correction

The second correction pass reduces the enhancement signal $\hat{d}$ to avoid artifacts such as overshoots and gradient inversions. This implements the third and fourth constraints of the optimization problem (11). The pixels are updated in a sequential order via two raster scans, one forward (left-to-right and top-to-bottom) and one backward. Approximating the partial derivatives by the finite difference kernel $[-1\ 1]$, and fixing the values of the pixel neighbors in (11) leads to the update equation

$$\begin{aligned}
&\text{If } \hat{d}_{ij} \geq 0, \\
&\begin{cases} \tau \leftarrow \min\limits_{(k,l)\in\mathcal{N}_4(i,j)} \Big(\Delta_{ij}^{kl} \left(\beta_{max}\mathbf{1}_{\Delta_{ij}^{kl}\geq 0} + \beta_{min}\mathbf{1}_{\Delta_{ij}^{kl}<0}\right) \\ \qquad\qquad + \varepsilon + \hat{c}_{kl} - \hat{c}_{ij} + \hat{d}_{kl} \Big), \\ \hat{d}_{ij} \leftarrow \max(0, \min(\tau, \hat{d}_{ij})), \end{cases} \\
&\text{otherwise,} \\
&\begin{cases} \tau \leftarrow \max\limits_{(k,l)\in\mathcal{N}_4(i,j)} \Big(\Delta_{ij}^{kl} \left(\beta_{min}\mathbf{1}_{\Delta_{ij}^{kl}\geq 0} + \beta_{max}\mathbf{1}_{\Delta_{ij}^{kl}<0}\right) \\ \qquad\qquad - \varepsilon + \hat{c}_{kl} - \hat{c}_{ij} + \hat{d}_{kl} \Big), \\ \hat{d}_{ij} \leftarrow \min(0, \max(\tau, \hat{d}_{ij})), \end{cases}
\end{aligned} \tag{15}$$

where $\Delta_{ij}^{kl} = \tilde{s}_{ij} - \tilde{s}_{kl}$, $\mathcal{N}_4(i,j)$ is the 4-neighborhood around pixel (i,j), and $\mathbf{1}_{(.)}$ is the zero-one function which takes value 1 when its subscript is true and 0 otherwise.

4. EXPERIMENTAL RESULTS

We present experimental results on two real images (spine and skull) over which synthetic patterns have been added to help study artifacts.

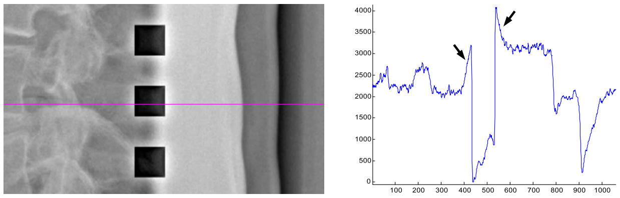

(a) Laplacian pyramid

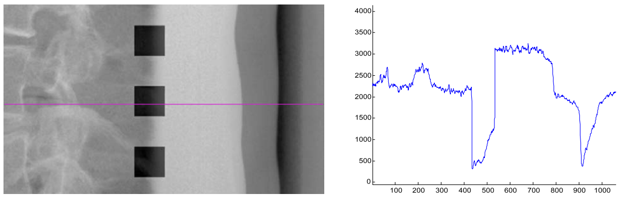

(b) Proposed algorithm

Fig. 4: A crop from the spine image enhanced by a classical Laplacian-based method (a), and by the proposed algorithm (b). The 1D intensity profiles are taken along the horizontal magenta line. The proposed algorithm avoids haloes around edges.

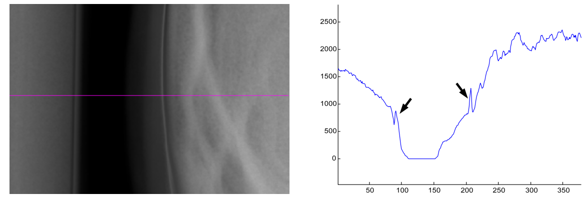

(a) Bilateral filter

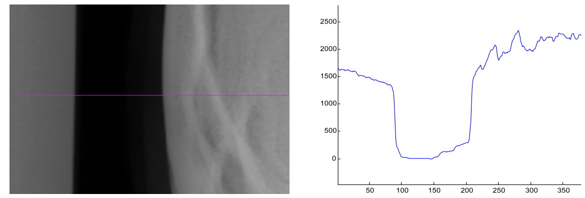

(b) Proposed algorithm

Fig. 5: A crop from the skull image enhanced using the bilateral filter (a), and by the proposed algorithm (b). The 1D intensity profiles are taken along the horizontal magenta line. The proposed algorithm avoids gradient inversions in slowly varying regions.

The images are 3114 by 3115 in size, for a total of about 9.7Mpx.

The software implementation reduces the image range by a factor of four, from 14b ($\max s = 16383$) to 12b ($\max \hat{s} = s_{max} = 4095$). The experiments have been run with a 6-level pyramid. The gains α have been set to 0 at the finest level to reduce the noise amplification and make the algorithm run faster, and to 3 at the other levels. The parameters σ, ε, β_{min}, and β_{max} have been set respectively to $10^{-1}s_{max}$, $10^{-3}s_{max}$, $3/4$ and 5. The experiments have been run on a Pentium 4 at 2.8GHz. In spite of the large image size, it takes only 1.95s to process each image.

Figure 3 shows a crop of the skull image after applying the global transfer function, and after applying weak-structure enhancement. It confirms the ability of the proposed algorithm to preserve the contrast of both weak and strong image structures.

Figure 4 shows a crop of the spine image enhanced without artifact correction (classical pyramid enhancement), and with the proposed method. Unlike classical pyramid enhancement, the proposed method does not generate haloes in regions surrounding large intensity variations, like those around the synthetic squares for instance.

Figure 5 shows a crop of the skull image enhanced with bilateral filtering [5], and with the proposed method. Unlike bilateral filtering, the proposed method does not create gradient inversions in slowly varying regions, like those surrounding the dark central area for instance.

5. CONCLUSION

In this paper, we have introduced a fast compression method based on the Laplacian pyramid, which avoids common artifacts such as loss of contrast, haloes and gradient inversions. We have shown that the artifact correction can be formulated as a linear-programming problem, for which we have proposed an efficient approximate solution. The effectiveness and speed of the proposed algorithm have been demonstrated on real medical images. Future work shall aim at generalizing the proposed method to image sequences.

References

[1] K. Devlin, A. Chalmers, A. Wilkie, and W. Purgathofer, "Tone reproduction and physically based spectral rendering," in *Eurographics*, 2002. 1

[2] G. W. Larson, H. Rushmeier, and C. Piatko, "A visibility matching tone reproduction operator for high dynamic range scenes," *IEEE Trans. on Visualization and Computer Graphics*, vol. 3, pp. 291–306, 1997. 1

[3] F. Drago, K. Myszkowski, T. Annen, and N. Chiba, "Adaptive logarithmic mapping for displaying high contrast scenes," *Computer Graphics Forum*, vol. 22, pp. 419–426, 2003. 1

[4] E. Reinhard, M. Stark, P. Shirley, and J. A. Ferwerda, "Photographic tone reproduction for digital images," in *SIGGRAPH*, 2002. 1

[5] F. Durand and J. Dorsey, "Fast bilateral filtering for the display of high-dynamic-range images," in *SIGGRAPH*, 2002. 1, 4

[6] Y. Li, L. Sharan, and E. H. Adelson, "Compressing and companding high dynamic range images with subband architectures," in *SIGGRAPH*, 2005. 1

[7] R. Fattal, D. Lischinski, and M. Werman, "Gradient domain high dynamic range compression," in *SIGGRAPH*, 2002. 1

[8] S. Bae, S. Paris, and F. Durand, "Two-scale tone management for photographic look," in *SIGGRAPH*, 2006. 1

[9] P. J. Burt and E. H. Adelson, "The Laplacian pyramid as a compact image code," *IEEE Trans. on Com.*, vol. COM-31, no. 4, pp. 532–540, 1983. 1

[10] D.G. Luenberger, *Linear and Nonlinear Programming*, Kluwer Academic Publishers, second edition, 2003. 1, 2, 3

[11] D. Taubman and M. Marcellin, *JPEG2000: Image Compression Fundamentals, Standards and Practice*, Springer-Verlag, 2001. 3