

JPEG2000 ROI Coding Through Component Priority for Digital Mammography

Joan Bartrina-Rapesta, Joan Serra-Sagristà, and Francesc Aulí-Llinàs

*The authors are with the Department of Information and Communications Engineering,
Universitat Autònoma de Barcelona, 08290 Cerdanyola del Vallès, Spain (e-mail:
jbartrina@deic.uab.es; joan.serra@uab.cat and fauli@deic.uab.es).*

Abstract

Region Of Interest (ROI) coding is a prominent feature of some image coding systems aimed to prioritize specific areas of the image through the construction of a codestream that, decoded at increasing bit-rates, recovers the ROI first and with higher quality than the rest of the image. JPEG2000 is a wavelet-based coding system that is supported in the Digital Imaging and Communications in Medicine (DICOM) standard. Among other features, JPEG2000 provides lossy-to-lossless compression and ROI coding, which are especially relevant to the medical community. But, due to JPEG2000 supported ROI coding methods that guarantee lossless coding are not designed to achieve a high degree of accuracy to prioritize ROIs, they have not been incorporated in the medical community.

This paper introduces a ROI coding method that is able to prioritize multiple ROIs at different priorities, guaranteeing lossy-to-lossless coding. The proposed ROI Coding Through Component Prioritization (ROITCOP) method uses techniques of rate-distortion optimization combined with a simple yet effective strategy of ROI allocation that employs the multi-component

support of JPEG2000 codestream. The main insight in ROITCOP is the allocation of each ROI to an component. Experimental results indicate that this ROI allocation strategy does not penalize coding performance whilst achieving an unprecedented degree of accuracy to delimit ROIs.

The proposed ROITCOP method maintains JPEG2000 compliance, thus easing its use in medical centers to share images. This paper analyzes in detail the use of ROITCOP to mammographies, where the ROIs are identified by computer-aided diagnosis. Extensive experimental tests using various ROI coding methods suggest that ROITCOP achieves enhanced coding performance.

Keywords:

Digital mammogram compression, JPEG2000 standard, Region Of Interest coding, Rate-Distortion optimization

1. Introduction

The American Cancer Society (ACS) [1] indicates that the probability of developing invasive breast cancer for USA's women younger than 39 is 1 in 210, and aged from 40 to 59 is 1 in 26. Mammography is the most effective way of detecting breast cancer before the onset of clinical symptoms [2].

The latest digital devices used in medical scenarios capture Mammograms and Angiograms images with a bit-depth resolution of 8-, 12- or 16-bits per pixel. In some cases, this high bit-depth resolution may produce files that grow to as much as 200 MB per mammography. Considering that current ACS guidelines for breast cancer screening recommend one annual mammography for women over 40 years of age, the increment in cost of both the

transmission and storage capacity for mammographies is raising every year. A medical center that produces 20 mammograms per day, for instance, requires storage capabilities of more than 4 GB per day [3], and of more than 1.4 TB per year. To manage such amount of data, the medical community uses Picture Archiving and Communication Systems (PACS) to store, retrieve, distribute, and display medical images. PACS systems are commonly constituted of large computer networks, servers, and workstations [4]. The Digital Imaging and Communications in Medicine (DICOM) standard [5] is used to store and distribute images in PACS.

Compression of medical images is commonly used in medical centers to reduce the amount of information needed to store images. Due to clinical needs, lossless compression of medical images is often a sensible choice [6]. There exist two approaches to losslessly encode images: *pure lossless*, which needs the decoding of the full codestream to recover the image, and *progressive lossy-to-lossless*, which allows a progressive decoding of the image and a better rate-distortion performance.

Among the currently available image compression standards, JPEG2000 [7] excels for its excellent coding performance and provision of advanced features. Since November 2001 it is included in DICOM, and it has been chosen in several medical centers to share and transmit medical images. Two of the most important features of JPEG2000 for the medical community are support for lossy-to-lossless compression, and Region Of Interest (ROI) coding. As stated previously, lossy-to-lossless is preferred over a pure lossless approach. ROI coding stands for the ability of the coding system to emphasize the most relevant areas of the image within the codestream so that when the image




Figure 1: ROIs defined by a CAD are emphasized in yellow.

is progressively transmitted and decoded those relevant areas are recovered first and with higher quality than the rest of the image.

On the other hand, due to recent advances in the field of image processing, every day it becomes more common that computer-aided diagnosis (CAD) assist radiologists in the detection of the most important features found in medical images. There exist several CAD algorithms that aid radiologists in the identification of pathologies on mammograms [8, 9, 10, 11], meniscal tears on magnetic resonance images [12, 13], and lung nodules on Computer Tomography [14, 15] or prostate cancer on Magnetic Resonance Imaging [16] images, among others. In the case of mammograms, CAD algorithms often delimit areas with sizes that vary from 0.002% to 4.5% of the mammogram size [17, 18] (see Figure 1).

In this context, it is a natural association to link CAD algorithms with

ROI coding. Radiologists are interested on the relevant areas needed to perform a correct diagnostic [19]. ROI coding provides enhanced quality for such ROIs. The benefits obtained from ROI coding are enormous. For example, Figure 2 depicts an area of a mammography determined by a CAD that is encoded using classic lossy-to-lossless compression, and lossy-to-lossless compression with ROI coding. This figure shows the decoded image in three instants of time when the codestream is transmitted over a network with a data rate of 100 Kbits/s. Only 8 seconds are required to recover the ROI at perfect quality when ROI coding is employed, whereas the classic approach requires 351 seconds (close to six minutes) to achieve the same quality. In scenarios such as telemedicine, or remote medical diagnosis, ROI coding is fundamental.

ROI coding methods have been widely studied in the literature. First approaches were proposed in mid-90s [20] and, more recently, there have appeared ROI coding methods for most state-of-the-art coding systems, such as Set Partitioning in Hierarchical Trees (SPIHT) [21], Set-Partitioning Embedded Block Coder SPECK [22], Three-dimensional Subband Block Hierarchical Partitioning (3D-SBHP) [23], or Progressive Resolution coding (PROGRES) [24], among others. Another trend to code ROIs of an image is to use techniques of shape adaptive coding based on SPIHT [25], or SPECK [26]. Though some of these methods achieve competitive coding performance, none of them is included in DICOM, which is a *conditio sine qua non* in medical centers.

The purpose of this research is to introduce a new JPEG2000 compliant ROI coding method providing progressive lossy-to-lossless coding that is able


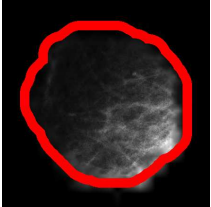

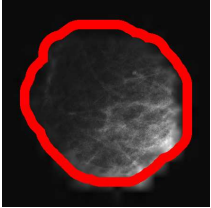

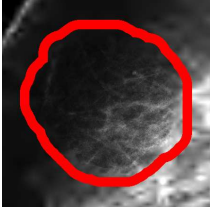
Classical Coding	ROI Coding
	
1 second	1 second
	
4 seconds	4 seconds
	
8 seconds	8 seconds

Figure 2: Visual comparison of a classical coding and a ROI coding method transmitted over the network at 100 kbits/s.

to exactly and losslessly recover the desired region, that allows to encode arbitrary shaped ROIs, that supports multiple ROI coding with different degrees of ROI priority, and that may deal with any bit-depth resolution of medical images. Our goal is to bring to the medical community the advantages of ROI coding.

This paper is structured as follows: Section 2 serves as a short review of

JPEG2000, and describes the mechanisms used in JPEG2000 for ROI coding; Section 3 introduces the proposed ROI coding method; Section 4 provides numerical and visual results assessing the performance of the proposed method when single and multiple ROIs are coded; and Section 5 concludes this work pointing out some remarks.

2. JPEG2000 Overview and ROI Coding

Most JPEG2000 [7] implementations require four main coding stages to produce a compliant codestream [27]: sample data transformations, sample data coding, rate-distortion optimization, and codestream re-organization. The main operations related to ROI coding in JPEG2000 are the fractional bitplane coding process carried out in sample data coding, and the rate-distortion optimization stage.

The JPEG2000's *fractional bitplane coder* is based on Embedded Block Coding with Optimized Truncation (EBCOT) [28]. The main idea behind this coding paradigm is to code small sets of wavelet coefficients (called codeblocks) independently, and to optimally truncate the bitstreams generated for these codeblocks to form the final codestream. If K_i denotes the number of bitplanes needed to represent all coefficients within codeblock \mathcal{B}_i , the fractional bitplane coder encodes each coefficient of \mathcal{B}_i from the highest bitplane $p = K_i - 1$ to the lowest bitplane $p = 0$. Each bitplane is coded using three coding passes that encode first that information supplying the greatest reductions in distortion. The bitstream generated for each codeblock can be truncated at the end of each coding pass, which produces several truncation points that can be potentially employed by the rate-distortion optimization

stage.

The aim of the *rate-distortion optimization* stage is to manage the bitrate and/or the distortion of the final codestream. When the user specifies a desired bitrate, the rate-distortion optimization stage maximizes the quality of the final codestream; when the desired quality is specified, the bitrate of the codestream is minimized. The Post Compression Rate-Distortion optimization (PCRD) [28] is the most common method to conduct this optimization process. PCRD uses the bitrate and the distortion of potential truncation points of bitstreams to pose the optimization problem through a generalized Lagrange multiplier for a discrete set of points. Let n_j denote the potential truncation points of the bitstream produced for codeblock \mathcal{B}_i , and let $R_i^{n_j}$ and $D_i^{n_j}$ denote, respectively, the bitrate and distortion of these points, with $R_i^{n_j} \leq R_i^{n_{j+1}}$. PCRD computes the rate-distortion slope $S_i^{n_j} = \Delta D_i^{n_j} / \Delta R_i^{n_j}$, with $\Delta D_i^{n_j} = D_i^{n_{j-1}} - D_i^{n_j}$ and $\Delta R_i^{n_j} = R_i^{n_j} - R_i^{n_{j-1}}$, identifying those truncation points with strictly decreasing rate-distortion slope or, in other words, those points lying on the convex hull. When Mean Squared Error (MSE) is the considered distortion metric, the distortion of each coding pass is determined as

$$D_i^{n_j} = G_{b_i} \sum_{k \in \mathcal{B}_i} (y[k] - \hat{y}^{n_j}[k])^2 \quad (1)$$

where $y[k]$ denotes the coefficients of codeblock \mathcal{B}_i , $\hat{y}^{n_j}[k]$ denotes quantized coefficients at truncation point n_j , and G_{b_i} stands for the energy gain factor of subband b_i to which codeblock \mathcal{B}_i belongs. Once the operational rate-distortion function of codeblocks is identified, PCRD selects those coding

passes with the highest rate-distortion slope values until the target bit-rate is attained.

2.1. ROI Coding Mechanisms

To prioritize a specific area of an image, JPEG2000 ROI coding methods identify first those samples belonging to the ROI in the wavelet domain, called ROI coefficients. Then, ROI coefficients are prioritized in order to recover them at higher quality than the rest of the image, the background. To carry out this prioritization process, JPEG2000 provides two main mechanisms: either modifying wavelet coefficients, or modifying distortion estimates carried out in the rate-distortion optimization stage, more precisely, in Equation (1).

ROI coding methods based on *modifying wavelet coefficients* take advantage of the fractional bitplane coder by means of the multiplication of ROI coefficients by a desired priority, say \mathcal{U} , in the wavelet domain. Through this multiplication the magnitude of ROI coefficients is higher than that of the background, thus the bitplane coder encodes first ROI coefficients. To speed up the prioritization process, \mathcal{U} is commonly chosen to be a power of 2, thus the multiplication is implemented as a bit-shift operation and is conceptually seen as a bitplane shift. The JPEG2000 standard supports two methods based on this mechanism: the MaxShift [27] and the Scaling [29] based method. The main difference between them is that the Scaling allows the user to choose \mathcal{U} , whereas in MaxShift \mathcal{U} is chosen to shift up all ROI coefficients above the background. For the Scaling method, the ROI shape must be rectangular or elliptic, and it is explicitly transmitted to the decoder, whereas MaxShift allows any ROI shape and it does not need to explicitly transmit it since it is implicitly coded within the bitstream. Several modi-

fications of these methods have been proposed in the literature, all of them aimed to more precisely combine the ROI with the background: Bitplane-by-Bitplane Shift [30] was the first method suggesting an interleaving of ROI and background bitplanes; [31, 32, 33] introduce modifications to the interleaving strategy to allow richer combinations of the ROI with the background. Though some of these methods provide better performance than those ones supported in JPEG2000, none of them kept JPEG2000 compliance. On the other hand, some contributions can be found in the literature aimed to compress mammograms combining techniques of fractal-based segmentation with JPEG2000 ROI coding [34, 35]. The only method based on the modification of wavelet coefficients able to maintain JPEG2000 compliance is introduced in [36], proposing a shift of ROI coefficients belonging to different subbands aimed to prioritize multiple ROIs with different degrees of interest.

On the other hand, methods based on the modification of distortion estimates increase the distortion estimation $D_i^{n_j}$ depending on whether the codeblock has ROI coefficients or not. If codeblock \mathcal{B}_i contains ROI coefficients, $D_i^{n_j}$ is modified according to

$$D_i^{n_j} = \mathcal{U} * D_i^{n_j} . \quad (2)$$

In this case, \mathcal{U} is not restricted to be a power of 2. Since D' is the considered measure by the Lagrange multiplier, codeblocks containing ROI coefficients are effectively more prioritized than codeblocks containing background coefficients. The drawback of this mechanism is that it is not able to discriminate ROI and background coefficients accurately since it only distinguishes ROI

and background codeblocks on a block by block basis. Thus, the performance of such mechanism may be severely degraded in some cases. The first ROI coding method based on this technique was the Implicit [7, Ch.16.2]. Sanchez, Basu, and Mandal [37] proposed a rearrangement of the codestream to allocate the ROI before the background, and ROI coding with separated codeblock was presented in [38]. [39] introduced the Subblock and Weighted ROI coding methods, which are able to better estimate the rate-distortion contributions of ROI codeblocks.

2.2. Features of JPEG2000 ROI coding mechanisms

Methods that modify wavelet coefficients provide excellent accuracy to delimit the ROI area (also referred herein to as fine-grain accuracy), however, they can not be generally applied in medical environments since the dynamic range supported by the decoder can be exceeded due to the high bit-depth resolution of medical images. For example, although the MaxShift method permits to encode ROIs with priority $\mathcal{U} \leq 37$ – which should be enough for medical images – compliant decoders are only compelled to decode K_{min} magnitude bitplanes, with a K_{min} value that depends on the decoder’s class [40, App.1] and that it is generally much lower than 37 due to the common use of data structures based on architectures of 32 bits. This may leave some background areas of the image with null coefficients, or with very poor quality, preventing the use of MaxShift and similar methods in medical centers, since loss of information may occur [27, H.3]. This shortcoming is referred herein to as *dynamic range problem*. All popular implementations of JPEG2000, such as Kakadu [41], Jasper [42], OpenJPEG [43] and JJ2000 [44], are *not* able to encode images of 16 bits per sample or more using the MaxShift




Figure 3: Bitplanes representation for original samples, and when Scaling and MaxShift ROI coding methods are performed.

ROI coding method due to the dynamic range problem. Figure 3 illustrates the dynamic range problem depicting the increase produced in the dynamic range of ROI coefficients (depicted in gray) after the modification of wavelet coefficients.

Even though methods based on the modification of distortion estimates overcome the dynamic range problem, they are not able to achieve the intended fine-grain accuracy to delimit ROIs. The motivation behind the ROI coding method proposed in this paper is to avoid the drawbacks of currently available ROI coding methods for medical images, on the one hand, the dynamic range problem raised by methods based on the modification of wavelet coefficients and, on the other hand, the poor fine-grain accuracy of rate-distortion based methods.

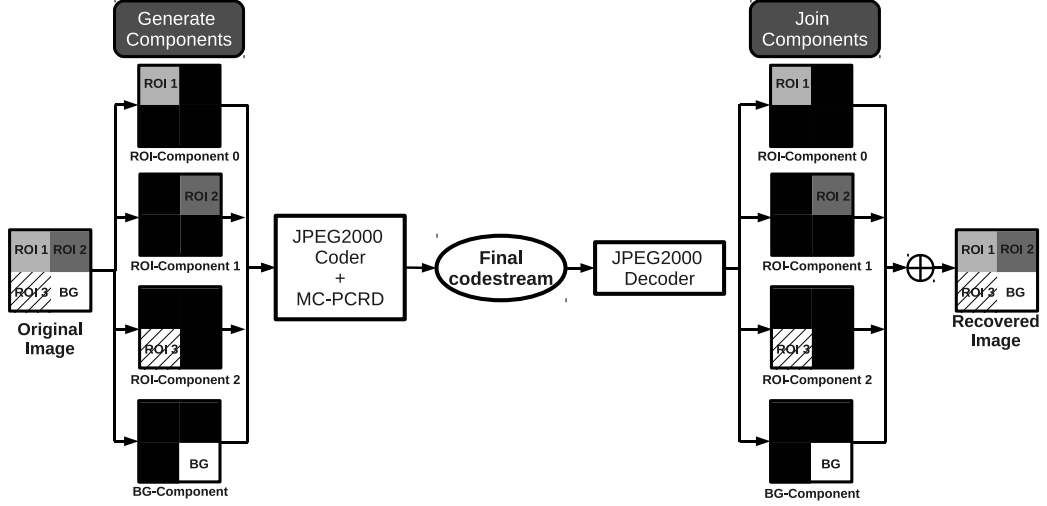


Figure 4: Operations for the proposed ROI coding method. Two operations are added in coder/decoder pipeline: Generate Components and Join Components.

3. ROI coding Through Component Priority

3.1. Main Insights

Two relevant features of JPEG2000 to our purposes are: support for multi-component images, and component scalability. Component scalability stands for the ability of the coding system to allow the access and manipulation of components in the compressed domain without needing to decompress the image. These features are employed by the proposed ROI coding Through Component Priority (ROITCOP) to allocate each ROI in an component where the non-ROI area is set to zero. Then, through the use of rate-distortion optimization techniques these components are prioritized at desired priorities, generating a multi-component image with each ROI prioritized at will.

To apply ROITCOP, the JPEG2000 core coding system requires two ad-

ditional operations in the coding pipeline, and a slight modification of the PCRD method. We note that the ROITCOP requires a JPEG2000 encoder that implements some rate distortion optimization method. Figure 4 depicts these two operations, called generate components, and join components. *Generate components* is an operation carried out in the encoder that defines as many components as ROIs have the image (referred to as ROI-components), plus one component for the background (referred to as BG-component). The operation *Join Components* sets the magnitude of each ROI coefficient to that recovered at the ROI-component with highest priority containing that ROI coefficient ¹. The magnitude of BG coefficients is set to that recovered at the BG-component.

For these Multi-Component (MC) images, a MC-PCRD is applied to combine the bitstreams from all components, minimizing the overall distortion. To correctly prioritize the desired ROI-components, a modification to MC-PCRD is introduced, updating the distortion estimates for specific codeblocks and components according to

$$D_{c,i}^{n_j} = \begin{cases} \mathcal{U}'_c * D_{c,i}^{n_j} & \text{if } c \in ROI \\ D_{c,i}^{n_j} & \text{otherwise} \end{cases}, \quad (3)$$

where $D_{c,i}^{n_j}$ denotes the distortion of component c for codeblock i at truncation point n_j , and \mathcal{U}'_c denotes the priority for ROI-component c . This operation modifies the distortion estimation for all ROI-components. Note

¹Note that a ROI coefficient may belong to more than one ROI in the wavelet domain due to the extension of coefficients.

that through Equation (3), rate-distortion estimates for coding passes are modified at codeblock level. The main difference with previous methods [7, Ch.16.2] [37, 39], is that in ROITCOP codeblocks contain only ROI or BG coefficients, whereas in previous approaches codeblocks may contain both types of coefficients. With the use of ROITCOP method, perfect fine-grain accuracy is achieved. Only two methods compliant with JPEG2000 are able to achieve comparable accuracy, MaxShift [27] and Scaling [29], though they suffer from the dynamic range problem.

3.2. ROITCOP codestream

As most ROI coding methods presented in the literature [7, 31, 45, 33, 36], the proposed ROITCOP method is able to prioritize several ROIs at once, with each ROI at a different priority. The maximum number of ROIs is $16384 - C$, where C is number of the components of the original image. This restriction is imposed by the maximum number of components allowed in JPEG2000, but it is large enough for most applications.

Though ROITCOP forces the encoding of the background area in ROI-components – which contain nothing –, the coding performance of the final codestream is almost not penalized since these areas are efficiently signaled within the codestream through headers. Headers within the codestream indicate which data is included for each codeblock. Practical experience indicates that the coding cost to signal headers for the components is negligible. For example, we have defined multiple ROIs for all images of the corpus presented in Section 4.1, and then each image has been encoded using ROITCOP to evaluate the extra information needed to encode components. Table 1 reports the size of the codestream when images are coded losslessly using ROITCOP,

Block size	Number of ROIs				
	0	1	2	3	4
32x32	5.65	+0.015	+0.027	+0.038	+0.049
64x64	5.63	+0.016	+0.028	+0.041	+0.052

Table 1: Average results for the image corpora used in the experiments when no ROIs and multiple ROIs are defined. Results report size of the compressed files, given in bits per sample.

compared to the classic compression without using ROI coding. Experiments are carried out using codeblock sizes of 32×32 , and 64×64 . Results suggest that the codestream size is penalized in no more than ~ 0.05 bits per sample (bps), which is negligible in practice.

3.3. Benefits of our proposal

To summarize, the key feature of our proposal is to allocate each ROI in a component and set coefficients of the non-ROI area of that component to zero. Then a modification of the distortion estimation is performed to prioritize the specific ROI within each component. Main advantages of this method are: 1) it avoids the dynamic range problem of the decoder; 2) it achieves very high fine-grain accuracy, comparable to that achieved by ROI coding methods based on modifying wavelet coefficients; 3) it is able to decode the ROI and the background in a lossy-to-lossless mode; 4) it enables the definition of multiple ROIs with different degrees of priority; 5) it is able to exclusively recover the desired ROI –simulating the MaxShift method–, through the component scalability, and; 6) it is compliant with JPEG2000 standard. Only a reassembly step is required to obtain an image from a Part

1.

4. Experimental Results

4.1. Materials and Parameters

The performance of ROITCOP is compared only to compliant JPEG2000 ROI coding methods using a set of conventional mammograms. ROI coding methods compliant with JPEG2000 based on coefficients modification are MaxShift and Scaling, while all ROI coding methods based on rate-distortion are compliant with JPEG2000.

Images are collected over a period of time from daily clinical activities at the Digital Medical Imaging Center, Parc Taulí Health Corporation [46] (Spain). Mammograms are acquired by the device TREX LORAD M-IV (Medical Jaco Equipment [47]), with a size of 2560×3328 , and a bit-depth of 12 bits per sample of signed data. Figure 5 depicts images and ROIs employed in this section. Rectangular and arbitrarily shaped ROIs are emphasized in green and red, respectively. ROI areas represent approximately 0.5% of the image area. Since the Scaling based method only allows rectangular and elliptical ROI shapes, comparisons for such method uses the ROI depicted in green in Figure 5, instead of the red one.

All methods are implemented in our JPEG2000 Part 1 implementation BOI [48]. Experiments are carried out using following coding parameters: 5 levels of 5/3 IWT (reversible), codeblock sizes of 32×32 or 64×64 , and restart coding variation ². Codeblocks of size 32×32 are recommended in [7,

² The restart coding variation [7, Chapter 12.4] is active, it maintains the coding passes

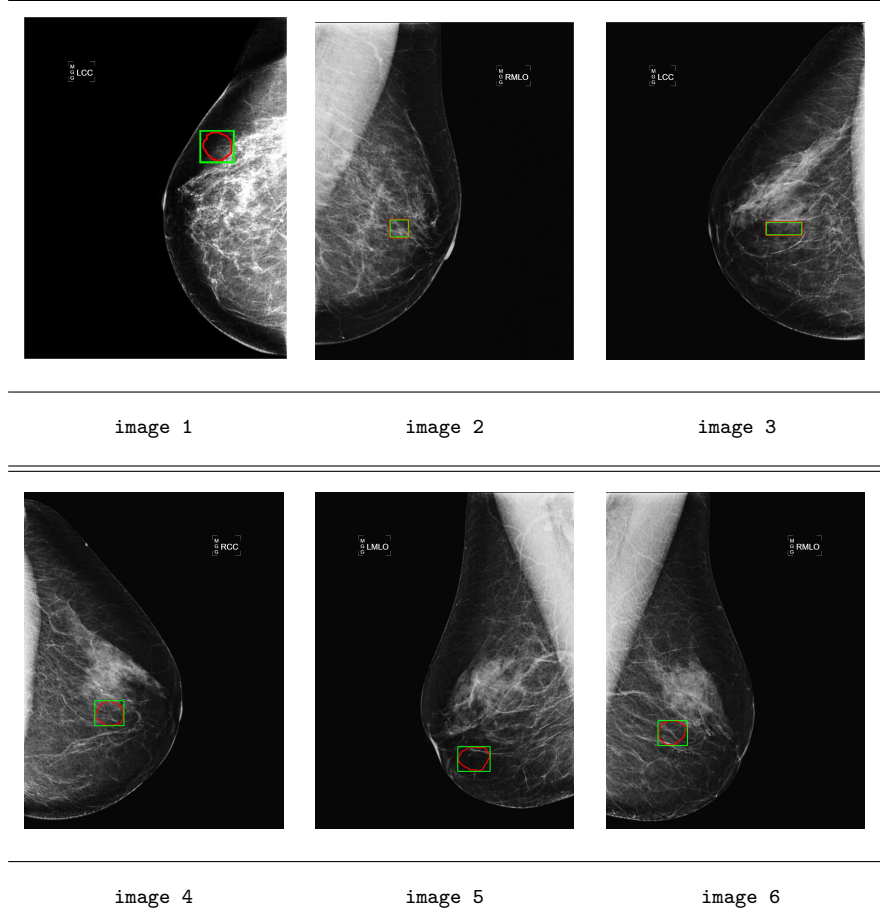


Figure 5: Image Set. Image size is 2560×3328 , bit-depth is 12 bits per sample.

Ch.16.2] to minimize the penalty of the Implicit method when distinguishing ROI codeblocks from background codeblocks.

Images are encoded at different target bitrates, decoded, and the image quality is assessed separately for the ROI and the background. The image

lengths in the codestream. It is devised to allow intra-codeblock parallelization for the bitplane coding stage, and for error resilience. It is used by default in BOI.

quality is reported in *Signal-to-Noise Ratio (SNR)*, which is a measure that accounts for the similarity between the original image I and the recovered image I^* , considering the signed data. SNR is defined as $10 \log_{10} \frac{\sigma^2}{MSE}$, where MSE is computed as $\frac{1}{N_x} \frac{1}{N_y} \sum_i \sum_j (I_{ij} - I_{ij}^*)^2$, and σ^2 denotes the variance of the original image. Higher SNR represents better quality for the recovered image I^* .

4.2. Numerical Results

Three scenarios are set up to compare ROITCOP with other compliant JPEG2000 ROI coding methods: 1) a single arbitrary ROI is defined and the priority is set to recover the ROI before the background (this test is intended to evaluate the fine-grain accuracy achieved by the evaluated methods); 2) arbitrary and rectangular ROIs are defined with priorities that combine ROIs and background (this test is aimed to illustrate the performance of evaluated methods when the ROI is interleaved with the background) and; 3) multiple ROIs with different priorities are defined (this test evaluates the order in which ROIs are recovered).

Test 1: Single arbitrary ROI

This test compares ROITCOP with MaxShift [27], Implicit [7, Ch.16.2], Subblock [39] and Scaling [29] methods. Figure 6 depicts the bitrate needed to recover the ROI area at a quality of 15 dB, 30 dB, 45 dB, or perfect recovery (lossless), on average for all images. The ROI priority is set equivalently for all methods.

ROITCOP and MaxShift recover the ROI at the same performance, suggesting that both methods achieve the same fine-grain accuracy, which is




Figure 6: Bits per sample (bps) needed to achieve a specific quality for the ROI and the background area. Bps is in average for all six images presented in Figure 6. The ROI priority is equivalent for all methods, and codeblock size is 32×32 . Bars indicate the bps needed to recover the ROIs and the background at specific qualities. Images have been encoded from 0.001 to 6 bps with a bit-rate step of 0.001.

slightly superior to that of the Scaling method. Note that MaxShift and Scaling may discard some of the least significant bitplanes of the background due to the dynamic range problem, producing loss of information. ROI coding methods that do not exceed the dynamic range (i.e., Implicit, and Subblock) guarantee lossless compression, however, they do not achieve the fine-grain accuracy achieved by MaxShift and ROITCOP. As an example, in a telemedicine scenario employing a 3G communication network, the ROITCOP method would require 5.83 seconds to recover the ROI, whereas the Implicit would need 21.66 seconds. The coding performance for the background is nearly equivalent for all methods and images.

Test 2: Fine-grain accuracy combining ROI and background

This test evaluates the ROI coding performance of ROITCOP compared to Scaling, Implicit, and Subblock methods, with different priorities, and for arbitrary and rectangular shaped ROIs. MaxShift is not included in these

experiments because it does not allow the combination between ROI and background coefficients.

Figure 7 (a), 7 (b), 7 (c) and 7 (d) provide the ROI coding performance, plotting the SNR, achieved by the arbitrarily defined ROIs (red ones in Figure 5), evaluated on two mammographies. ROITCOP outperforms Scaling, Subblock and Implicit methods in all cases. The background coding performance is nearly the same for all methods.

Figure 7 (e) and 7 (f) depict results obtained when rectangular ROIs are defined in one mammography. Results suggest that ROITCOP and Scaling methods achieve the best results.

It is also worth noting that when the codeblock size is set to 64×64 , the performance of ROITCOP and Scaling methods is maintained, while the performance of Subblock and Implicit methods degrades due to the rough discrimination between blocks that belong to the ROI or to the background.

Test 3: Multiple ROIs

The evaluation of the ROI coding performance for multiple arbitrary ROIs with different priorities is carried out comparing ROITCOP and Subblock. Implicit method is discarded because it does not achieve the intended fine-grain accuracy, MaxShift and Scaling methods are not considered since they can not assure lossless recovery of the ROI areas, and MaxShift can not prioritize multiple ROIs with different priorities. The ROI performance of ROITCOP and Subblock is studied when a progressive lossy-to-lossless and lossless compression is used. Figure 8 depicts **image 1** and **image 3** with the ROIs identified for these experiments.

Figure 9 depicts the coding performance achieved by ROITCOP (solid

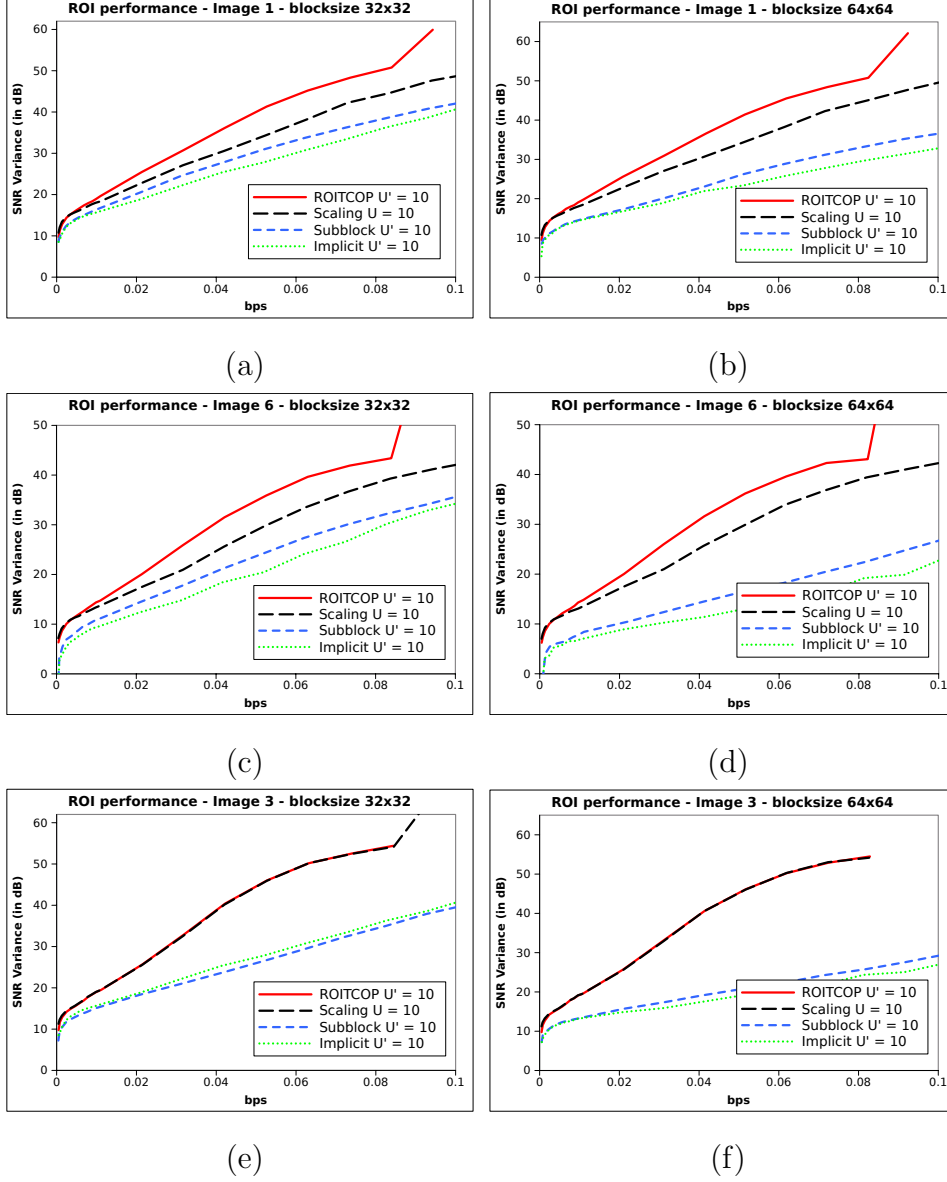


Figure 7: ROI coding performance for the arbitrary ROI of **image 1** and **image 6**, and rectangular ROI of **image 3**, showing the SNR of the ROI area achieved by ROITCOP, Scaling, Subblock and Implicit. Results are for a ROI priority $U' = 10$, for codeblock sizes set to 32×32 and 64×64 for progressive lossy-to-lossless compression.




Figure 8: Original images with multiple ROI definitions, (a) image 1, and (b) image 3.

line) and Subblock (dashed line) methods when three ROIs are defined and encoded at different priorities using codeblock sizes of 64×64 in a progressive lossy-to-lossless compression. The priorities for the different ROIs are set to: ROI_1 = 10, ROI_2 = 8, and ROI_3 = 6. Results suggest that our proposal works properly when multiple ROIs are defined, since ROI_1 has better quality than ROI_2 and ROI_3, and ROI_2 has better quality than ROI_3. In addition, for most of the bit-rates and ROIs, ROITCOP outperforms the coding performance of Subblock.

Figure 10 reports the bps needed for ROITCOP and Subblock methods to recover the three defined ROI areas and the whole image losslessly. When multiple ROIs are defined, ROITCOP needs, approximately, from 2 to 5 times less bitrate than Subblock to recover the ROIs.

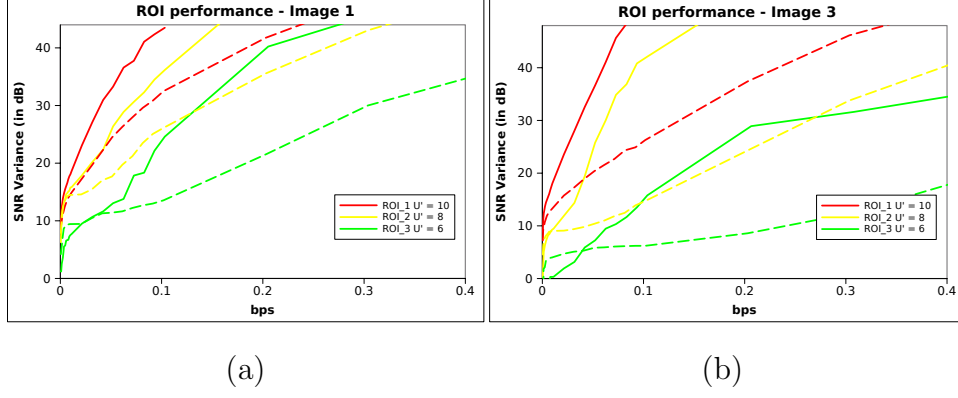


Figure 9: ROITCOP ROI coding performance for multiple ROIs prioritized with distinct weights, (a) image 1, and (b) image 3.

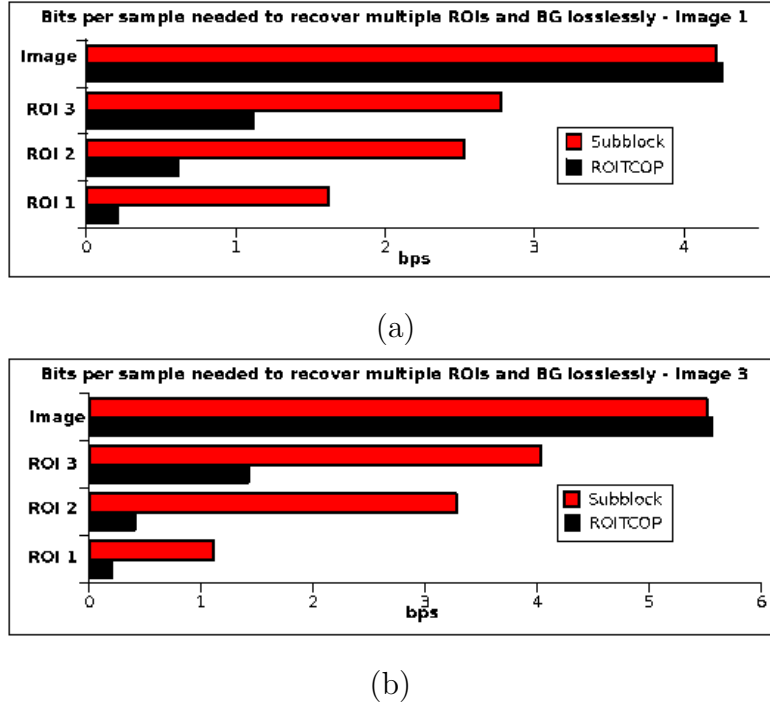


Figure 10: ROITCOP and Subblock coding performance for multiple ROIs, indicating the bps needed to recover the different ROIs and the whole image losslessly, (a) image 1, and (b) image 3.

4.3. Visual Comparison

To assess the performance of the ROI coding methods visually, MaxShift, Subblock, and ROITCOP are used to code an arbitrary shaped ROI. MaxShift is selected because it provides the best fine-grain accuracy of methods based on the modification of wavelet coefficients; whereas Subblock is chosen because it provides the best fine-grain accuracy of methods based on the modification of distortion estimates.

Figure 11 depicts the recovered images when the ROI area covers 0.6% of the image. ROITCOP recovers the ROI with the highest accuracy of all methods, helping the specialist to pay attention only to the relevant area. MaxShift does not delimit the edges of the ROI accurately, and Subblock recovers too much surrounding area around the ROI.

5. Conclusions

ROI coding is a particularly suitable coding mechanism for medicine activities, providing the possibility to adequately compress those regions of the image that have the highest diagnostic relevance. ROI coding allows to transmit these relevant areas earlier and at a higher quality than the rest of the image. In a tele-diagnosis scenario, progressive lossy-to-lossless compression is also desirable, therefore, JPEG2000 image coding standard is convenient since it is included in DICOM, supports ROI coding, and provides progressive lossy-to-lossless performance.

JPEG2000 standard supports ROI coding through two mechanisms: either modifying wavelet coefficients, or employing rate-distortion optimization techniques. Compliant JPEG2000 methods that modify wavelet coefficients




Figure 11: Visual comparison of a ROI in a digital mammogram. Figure shows (a) the Original image with ROI emphasized in red, and different ROI coding methods: (b) MaxShift, (c) ROITCOP, and (d) Subblock, at 0.01 bps, for a codeblock size of 64×64 . All images are scaled at 70% of its original size.

are MaxShift and Scaling. Both methods achieve high fine-grain accuracy, however, they suffer from the dynamic range problem, which may not guarantee a perfect lossless recovery for the ROI and the background. Methods that employ rate-distortion optimization techniques are the Implicit and Subblock. They ensure the recovering of the ROI and the background losslessly, but they recover too much information outside the ROI area.

This paper introduces ROITCOP ROI coding method. This method allocates each ROI to an component and uses the rate-distortion optimization techniques to prioritize the desired ROI. ROITCOP allows progressive lossy-to-lossless recovery for the ROI and the background, does not penalize the coding efficiency, permits to recover exclusively a requested ROI through the component scalability feature of JPEG2000, and achieves a fine-grain accuracy. ROITCOP needs a negligible extra bitrate to encode the additional headers coded in each component. Experimental results suggest that the ROI coding performance of ROITCOP outperforms Scaling, Implicit, and Subblock ROI coding methods, recovering the ROI areas at higher quality using equivalent parameters, while for MaxShift nearly the same ROI coding performance is obtained.

Applications that can benefit from the proposed ROITCOP method are those that need to: store the ROI and the background image losslessly due to legal issues and clinical needs; determine several ROIs with different degrees of importance, as defined by CAD or radiologists; and retrieve the diagnostic area at high levels of quality at low bitrates, for instance for mobile devices.

6. Acknowledgment

This work has been partially supported by the Spanish and Catalan Governments, and by FEDER, under Grants TIN2009-14426-C01, SGR2009-1224, 2008-BPB-0010, and RYC-2010-05671.

- [1] American Cancer Society, ACS (Sep. 2009).
URL <http://www.cancer.org>
- [2] H.-P. Chan, L. T. Niklason, D. M. Ikeda, K. L. Lam, D. D. Adler, Digitization requirements in mammography: Effects on computer-aided detection of microcalcifications, *AAPM Medical Physics* 21 (7) (1994) 1203–1211.
- [3] B. W. Orenstein, Eight keys to digital mammography success, *Radiology Today* 7 (19) (2006) 30 pages.
- [4] X. Cao, H. Huang, Current status and future advances of digital radiography and PACS, *IEEE Eng. in Medicine and Biology Magazine* 19 (5) (2000) 80–88.
- [5] Digital Image and Communication in Medicine, DICOM (Sep. 2009).
URL <http://medical.nema.org/>
- [6] V. Sanchez, R. Abugharbieh, P. Nasiopoulos, Symmetry-Based Scalable Lossless Compression of 3D Medical Image Data, *IEEE Trans. on Medical Imaging* 28 (7) (2009) 1062–1072.

- [7] D. Taubman, M. W. Marcellin, JPEG2000 Image compression fundamentals, standards and practice, Kluwer Academic Publishers, Norwell, Massachusetts 02061 USA, 2002.
- [8] Y. Hatanka, T. Hara, H. Fujita, S. Kasai, T. Endo, T. Iwase, Development of an automated method for detecting mammographic masses with a partial loss of region, IEEE Trans. on Medical Imaging 20 (12) (2001) 1209–1214.
- [9] M. Lado, P. Tahoces, A. Mendez, M. Souto, J. Vidal, Evaluation of an automated wavelet-based system dedicated to the detection of clustered microcalcifications in digital mammograms, Informatics for Health and Social Care 26 (3) (2001) 149–163.
- [10] C. Varela, P. G. Tahoces, A. J. Méndez, M. Souto, J. J. Vidal, Computerized detection of breast masses in digitized mammograms, Comput. Biol. Med. 37 (2) (2007) 214–226.
- [11] A. Cao, Q. Song, X. Yang, Robust information clustering incorporating spatial information for breast mass detection in digitized mammograms, ELSEVIER Computer Vision and Image Understanding 109 (1) (2008) 86–96.
- [12] N. Safdar, B. Ramakrishna, G. Saiprasad, K. Siddiqui, E. Siegel, Role of computer aided detection (CAD) integration: case study with meniscal and articular cartilage CAD applications, in: Proc. of SPIE PACS and Imaging Informatics, Vol. 6919, 2008, pp. 69190J–69190J–7.

- [13] B. Ramakrishna, W. Liu, G. Saiprasad, N. Safdar, C. Chang, K. Siddiqui, W. Kim, E. Siegel, J.-W. Chai, C. C.-C. Chen, S. K. Lee, An automatic computer-aided detection system for meniscal tears on magnetic resonance images, *IEEE Trans. on Medical Imaging* 28 (8) (2009) 1308–1316.
- [14] N. Niki, Y. Kawata, M. Kubo, A CAD system for lung cancer based on CT image, *ELSEVIER Computer Assisted Radiology and Surgery* 1230 (2) (2001) 631–638.
- [15] A. El-Baz, G. Gimel'farb, R. Falk, M. El-Ghar, A New CAD System for Early Diagnosis of Detected Lung Nodules, in: *Proc. of IEEE International Conference on Image Processing*, Vol. 2, 2007, pp. 461–464.
- [16] P. Puech, N. Betrouni, R. Viard, A. Villers, X. Leroy, L. Lemaitre, Prostate cancer computer-assisted diagnosis software using dynamic contrast-enhanced MRI, in: *Conference of the IEEE EMBS*, Vol. 1, 2007, pp. 5567–5570.
- [17] F. J. Gilbert, S. M. Astley, C. R. Boggis, M. A. McGee, P. M. Griffiths, S. W. Duffy, O. F. Agbaje, M. G. Gillan, M. Wilson, A. K. Jain, N. Barr, U. M. Beetles, M. A. Griffiths, J. Johnson, R. M. Roberts, H. E. Deans, K. A. Duncan, G. Iyengar, Variable size computer-aided detection prompts and mammography film reader decisions, *Breast Cancer Research* 10 (4) (2008) 8 pages.
- [18] U. Bottigli, P. Cerello, P. Delogu, M. Fantacci, F. Fauci, B. Golosio, A. Lauria, E. L. Torres, R. Magro, G. Masala, P. Oliva, R. Palmiero,

- G. Raso, A. Retico, S. Stumbo, S. Tangaro, GPCLAMA, a mammographic CAD in grid connection, in: Proc. of International Congress Series, Vol. 1, 2003, pp. 944–949.
- [19] C. Doukas, I. Maglogiannis, Region of interest coding techniques for medical image compression, *IEEE Engineering in Medicine and Biology Magazine* 26 (5) (2007) 29–35.
- [20] A. T. Duchowski, B. H. McCormick, Simple multiresolution approach for representing multiple regions of interest (ROIs), in: Proc. of Visual Communications and Image Processing, Vol. 2501, 1995, pp. 175–186.
- [21] S. Tasdoken, A. Cuhadar, ROI coding with integer wavelet transforms and unbalanced spatial orientation trees, in: Proc. of IEEE International Conference on Engineering in Medicine and Biology, Vol. 1, 2003, pp. 841–844.
- [22] Z. Li-Bao, Y. Xian-chuan, W. Shi-hong, New region of interest image coding based on multiple bitplanes up-down shift using improved SPECK algorithm, in: Proc. of IEEE International Conference on Innovative Computing, Information and Control, Vol. 3, 2006, pp. 629–632.
- [23] Y. Liu, W. Pearlman, Region of interest access with three-dimensional SBHP algorithm, in: Visual Communications and Image Processing, Vol. 6077, 2006, pp. 17–19.
- [24] C. Yushin, W. Pearlman, Hierarchical Dynamic Range Coding of Wavelet Subbands for Fast and Efficient Image Decompression, *IEEE Trans. Image Processing* 16 (8) (2007) 2005–2015.

- [25] A. Said, W. Pearlman, A new, fast and efficient image codec based on set hierarchical trees, *IEEE Trans. Circuits and Syst. Video Technol.* 6 (1996) 243–250.
- [26] G. Xie, H. Shen, A highly scalable SPECK image coder, in: *Proc. of IEEE International Conference on Image Processing*, Vol. 2, 2004, pp. 1297–1300.
- [27] ISO/IEC, Information technology - JPEG 2000 image coding system - Part 1: Core coding system (Dec. 2000).
- [28] D. Taubman, High performance scalable image compression with EBCOT, *IEEE Trans. Image Processing* 9 (7) (2000) 1158–1170.
- [29] ISO/IEC, Information technology - JPEG 2000 image coding system - Part 2: Extensions (Apr. 2004).
- [30] Z. Wang, A. C. Bovik, Bitplane-by-bitplane shift (BbBShift)- a suggestion for JPEG2000 region of interest image coding, *IEEE Signal Processing Letters* 9 (5) (2002) 160–162.
- [31] L. Liu, G. Fan, A new JPEG2000 region-of-interest image coding method: Partial significant bitplanes shift, *IEEE Signal Processing Letters* 10 (2003) 157–163.
- [32] L. bao Zhang, X. chuan Yu, New region of interest image coding using general layered bitplane shift for medical image compression, *International Journal of Computational Intelligence Research* 3 (1) (2007) 8 pages.

- [33] Z. Li-Bao, New approach to JPEG2000 region of interest image coding and its applications for remote sensing image, in: Proc. of SPIE, Geospatial Information, Data Mining, and Applications, Vol. 6045, 2005, pp. 952–956.
- [34] H.-Y. Chan, H. Sari-Sarraf, B. I. Grinstead, S. S. Gleason, Content-based compression of mammograms with fractal-based segmentation and a modified JPEG2000, SPIE Optical Engineering 3 (2986).
- [35] C.-S. Lo, P.-C. Chung, S.-K. Lee, G.-C. Hsu, Fractal based JPEG2000 Roi Coding, in: Proc. of IEEE International Conference on Systems, Man, and Cybernetics, Vol. 43, 2006, pp. 553–555.
- [36] P. G. Tahoces, J. R. Varela, M. J. Lado, M. Souto, Image compression: Maxshift ROI encoding options in JPEG2000, ELSEVIER Computer Vision and Image Understanding 109 (2) (2008) 139 – 145.
- [37] V. Sanchez, A. Basu, M. Mandal, Prioritized region of interest coding in JPEG2000, IEEE Trans. Circuits Syst. Video Technol. 14 (9) (2004) 1149–1155.
- [38] Y. Xie, G.-Q. Han, ROI coding with separated code block, in: Proc. of IEEE International Conference on Machine Learning and Cybernetics, Vol. 8, 2005, pp. 5447–5451.
- [39] J. Bartrina-Rapesta, J. Serra-Sagristà, F. Aulí-Llinàs, JPEG2000 ROI Coding with Fine-Grain Accuracy through Rate-Distortion Optimization Techniques, IEEE Signal Processing Letters 16 (1) (2009) 45–49.

- [40] ISO/IEC, Information technology - JPEG 2000 image coding system - Part 4: Conformance testing (May 2002).
- [41] D. Taubman, Kakadu Software Home Page (Sep. 2009).
URL <http://www.kakadusoftware.com/>
- [42] Adams, M.D., JasPer Project Home Page (Sep. 2009).
URL <http://www.ece.uvic.ca/~mdadams/jasper/>
- [43] OpenJpeg, OpenJpeg Home Page (Sep. 2009).
URL <http://www.openjpeg.org/>
- [44] EPFL and Canon Research France and Ericsson, Java implementation of JPEG 2000 (Sep. 2009).
URL <http://jpeg2000.epfl.ch>
- [45] Z. Li-Bao, W. Ke, New approach for JPEG2000 region of interest image coding hybrid bitplane shift, in: Proc. of IEEE International Conference on Machine Learning and Cybernetics, Vol. 6, 2004, pp. 3955–3960.
- [46] Coorporacio Parc Tauli, Sabadell, Spain (Sep. 2009).
URL <http://www.cspt.es>
- [47] Medical Jaco Equipment, JACO (Sep. 2009).
URL <http://www.jacomed.com/>
- [48] Group on Interactive Coding of Images, BOI software (Sep. 2009).
URL <http://www.gici.uab.es/BOI>

7. Biographies



Dr. Joan Bartrina-Rapesta

received the B.Sc., B.E., M.S. and Ph.D. degrees in Computer Engineering in 2002, 2004, 2006, and 2009, respectively, all from Universitat Autònoma de Barcelona. He was awarded with a doctoral fellowship from Universitat Autònoma de Barcelona. He has collaborated in the development of BOI, a JPEG2000 Part 1 implementation, and also in the development of CADI, a JPEG2000 Part 9 implementation. His research interests include a wide range of image coding topics, including highly scalable image and video coding systems, region-of-interest coding, rate-distortion optimization techniques, distortion estimation, and interactive image and video transmission.



Dr. Joan Serra-Sagristà,

Ph.D.'99 in Computer Science from Universitat Autònoma Barcelona (UAB), Spain, is an Associate Professor and Director of the Group on Interactive Coding of Images at UAB. From 1997 to 1998, and in 2000, he hold a DAAD research grant at University Bonn, Germany. His research interests include image coding, data compression, and signal processing, with applications to

remote sensing and telemedicine, having coauthored several papers. He is a 2006 recipient of the Intensification Program Young Investigator Award, and has served on the steering and technical program committees of several international conferences and major international journals.



Dr. Francesc Aulí-Llinàs

received the Ph.D. degree in Computer Engineering in 2006 from Universitat Autònoma de Barcelona. Since 2002 he has been awarded with doctoral and postdoctoral fellowships from the Spanish and Catalan Governments. From 2007 to 2009 he carried out two research stages of one year each with professors David Taubman and Michael Marcellin. He is the main developer of BOI, a JPEG2000 Part 1 implementation that was awarded with a free software mention from the Catalan Government, and serves on the technical program committees of several international conferences and journals.