

Selective coding with controlled quality decay for 2D and 3D images in a JPEG2000 framework

Alberto Signoroni, Fabio Lazzaroni, Riccardo Leonardi

Signals and Communications Lab. - DEA - University of Brescia, Italy
via Branze, 38 - I25123 Brescia, Italy Tel.+39030 3715432 Fax.+39030 380014

ABSTRACT

This paper presents some ideas which extend the functionality and the application fields of a spatially selective coding within a JPEG2000 framework. At first, the image quality drop between the Regions of Interest (ROI) and the background (BG) is considered. In a conventional approach, the reconstructed image quality steeply drops along the ROI boundary; however, this effect could be considered or perceived objectionable in some use cases. A simple quality decay management is proposed here, which makes use of concentric ROI with different scaling factors. This allows the technique to be perfectly consistent with the JPEG2000 part 2 ROI definition and description. Another considered issue is the extension of the selective ROI coding to a 3D Volume of Interest coding. This extension is currently under consideration for the part 10 of JPEG2000, JP3D. An easy and effective 2D to 3D extension for the VOI definition and description is proposed here: a VOI is defined by a set composition of ROI generated solids, where ROI are defined along one or more volume cutting direction, and is described by the relative set of ROI parameters. Moreover, the quality decay management can be applied to this extension. The proposed techniques could have a significant impact on the selective coding of medical images and volumes. Image quality issues are very important but very critical factors in that field, which also constitutes the dominant market for 3D applications. Therefore, some experiments are presented on medical images and volumes in order evaluate the benefits of the proposed approaches in terms of diagnostic quality improvement with respect to a conventional ROI coding usage.

Keywords: Selective ROI coding, image quality, wavelet coding, 3D datasets, JPEG2000, medical imaging

1. INTRODUCTION

In the last years, Region-of-Interest (ROI) or selective wavelet coding has achieved a great value due to an increasing importance of two different but correlated aspects: the scalability need (to meet heterogeneous communication and storage resources and Quality of Service issues) and the content analysis and indexing need (to manage and describe, at various semantic levels, the information contained by the visual data). Selective wavelet coding can be used to produce separate ROI and background (BG) coded streams or to produce a (scalable or progressive) bit-stream where the ROI quality is constantly maintained superior with respect to the BG one, until the lossless or near-lossless ROI representation is reached. The former approach usually exploits the possibility to have a wavelet representation of arbitrarily shaped regions; the latter realizes a coefficient ranking by using a subband coefficient weighting mask prior to the quantization stage. Irrespective of the ROI shape in the spatial domain, the subband weighting mask can be easily created by considering the coefficients involved in the wavelet (filtering and subsampling) decomposition of the ROI.^{1,2} The quality difference between ROI and BG can be modulated by deciding the value of the mask in the various subbands. One of the peculiarities of the still image coding standard JPEG2000 is the possibility to implement a ROI coding³ using at least two different approaches, the “max-shift” and the “scaling-based” methods, which are respectively described in the part 1 (baseline⁴) and part 2 (extensions⁵) of the standard. Despite those possibilities, the actual development of the ROI coding depends on some application dependent issues. In fact, there are plenty of potential selective coding application fields, but they usually require original solutions and a customized definition of the ROI coding modalities and communication scenarios. This significantly influences the choice of the appropriate ROI coding technique, but also the nature and degree of interaction with the user, the

Authors e-mail: {Alberto.Signoroni,Fabio.Lazzaroni,Riccardo.Leonardi}@ing.unibs.it

communication modalities, the storage policy and the application interface usability. Some of those issues are addressed in the part 9 of JPEG2000, JPIP,⁶ but the ROI usage remains critical when the preservation of the semantic meaning of the coded information is essential. Ensuring a semantically lossless coding while maximizing the compression ratio is a difficult task and represents a still open scientific question. The case of biomedical image compression is representative of the above issues, which risk to generate some refusing thoughts about the lossy coding usage, with the consequence of discarding the interesting opportunity of preserving the image semantic content by means of an adequate ROI coding. In addition, physician are not image processing experts while usually signal processing professionals are unaware of some medical image peculiarities. This generated confusion from a long time, but reliable positions, even if sporadic, begin to come out.⁷ On the other hand, it is interesting to note that a main motivation to the inclusion of ROI coding in the JPEG2000 standard just came from the medical imaging community. This is because of the increasing trend of the digital diagnostic imaging production, the necessity to communicate and store a great amount of data and the possibility to exploit the presence of Region of Diagnostic Interest into the images.⁸

Our contribution here is the introduction of some new techniques which make the ROI coding more flexible in the above perspective of having the possibility to cope with the open questions of the use of ROI coding in critical environments. Our additional attention is to introduce these aspects in consistency with the existing or the ongoing JPEG2000 standard definitions, with the endeavor to couple simple ideas with a high potential benefit for applications. We consider the conventional ROI co-decoding with mask generation in the wavelet domain, which is also described in the part 2 of JPEG2000 with the name of “scaling-based” method. In this framework we have the possibility to define one or more ROI and to qualify them with a scaling parameter, which is linked to the relative importance or “saliency” of the ROI with respect the rest of the image (the BG). In the case of overlapping multiple ROI (both in the image or wavelet domain) we assume, as it is reasonable and also foreseen in,⁵ that the maximum value among the intersecting ROI’s must be chosen (“max-rule”). In previous works^{9,10} we presented, especially for biomedical applications, the possibility and the need to qualify the ROI with another parameter, that cannot be directly handled within the considered framework, that is the degree of “focalization” of the ROI, which virtually determine the extent of a smoothed quality decay between ROI and BG.¹¹ In the case of medical imaging, we observed that a sharp quality disparity on the ROI boundaries could be perceived as an artifact itself that, depending on the diagnostic situation and the dimension of the ROI, could impair the “diagnostic quality” of the coded image. Thus, we proposed a masking function family with a sharp to smooth parameterized decaying profile. In this work we first propose an alternative (but virtually equivalent) way to manage this problem which is based on a set of concentric ROI’s with similar shape (rectangular or elliptical) and different dimension (Sec.2). Due to the max-rule for overlapping ROI’s, this set is equivalent to a staircase pyramid masking function, which inclination can be modulated, with a proper interaction, by a “focalization” parameter.

Another issue we want do deal with, is the possibility to easily define Volumes of Interest (VOI) for still 3D datasets (Sec.3). This is a topic of interest for the part 10 of JPEG2000, called JP3D.¹² The max-shift technique could be extended to run with free-hand drawn surfaces, however an accurate definition of a VOI is a time-consuming activity if made on single slices and involves considerable interaction problems if directly implemented in 3D. On the other hand, the scaling-based approach needs an explicit representation of the VOI, which should be compact but sufficiently flexible and representative of the interesting structures. We propose a VOI definition technique which is based on the possibility to slice the volume along various directions (typically the perpendicular axes) hence producing images on which define some ROI’s. That ROI’s will became the base of rectangular or elliptical cylindroids, the intersection (or eventually the union) of which quickly locates a 3D region. That region surrounds, for example, a 3D object of diagnostic interest. Thus we guarantee an effective definition and representation of VOI’s which make simple and compact the analytical 3D surface description and parameterization. This could be an interesting approach to the VOI definition in JP3D, which is conceived as an extension of part 1 and part 2 of JPEG2000. Moreover, the quality decay managing by concentric ROI can be straightforwardly extended to the volumetric case. In Sec.4 medical image applications are considered. Some examples of ROI and VOI coding are presented where a neuroradiologist evaluates the “diagnostic quality” of the compressed images with and without using the “focalization” parameter to qualify the ROI selection.

2. ROI CODING WITH COEFFICIENT MASKING AND PROGRESSIVE QUALITY DECAY

2.1. ROI coding and visual applications

As mentioned above, in the JPEG2000 standard there are two different region of interest (ROI) coding modalities. The first one is described in the “core coding system”,⁴ it uses a weighting mask to separate the ROI/BG coefficient range values and thus it provides for a separate ROI/BG coding rate with the benefit to not having to transmit the ROI shape. The second modality is described into the “extensions”⁵ and allows an integrated representation and transmission of the ROI and BG by maintaining a superior reconstruction quality into the ROI, with a scaling factor that can be decided by the user. In this case the ROI shape must be explicitly transmitted before the coded coefficients and must be a rectangle or an ellipsis, due to their synthetic parameterization which does not influence the coding performance, even with multiple ROI’s. Both these modalities determine a clear quality disparity in the correspondence of the ROI boundaries. The sharpness of this quality transition is only slightly mitigated by the action of the reconstruction filters.¹ In some situations, the user could perceive this transition as an artifact itself,^{9,10} therefore the need of a smoothed quality decay arises. This is the case of many situations in medical imaging, as we will see in Sec.4. Moreover, some studies^{13,14} argues that, from a perceptual point of view, an ideal ROI coding should follow the spatial resolution of the human visual system: the ROI corresponds to the foveation area and the surrounding reconstruction quality must be at least proportional to the photoreceptors density into the retina. The problem of those studies is the dynamic nature of the focal points, i.e. of the perceptual interest. At the same time, the semantic interest has a more static nature, especially when images are observed with specific objectives, but equally benefits from a perceptual optimization.¹⁰ In the following subsection the possibility to obtain a ROI quality decay, which might be modulated by the user with a perfectly conformable JPEG2000 part 2 technique, is explained. A set of concentric ROI which encircles the selected one are defined with a decreasing scaling factor.

2.2. Quality decay control with concentric ROI masks

In the “scaling-based” approach the ROI can be rectangular or elliptical. Both elliptical and rectangular ROI selections can be described by the coordinates (x_t, y_t) and (x_b, y_b) of the opposite top-left and bottom-right angles of the selected rectangle (which circumscribe the ROI in the case of elliptical selection); this is shown in Fig.1. Therefore, in the image domain, the ROI shape can be described by the vector $Ge = (x_t, y_t, x_b, y_b, t)$ (where t determines the rectangular or elliptical shape), and represented by a binary mask $ROI(Ge, x, y)$ as shown in Fig.2. We also define $a = \min(|x_t - x_b|, |y_t - y_b|)$ and $b = \max(|x_t - x_b|, |y_t - y_b|)$ as the smaller and greater diameter (side) of the ellipsis (rectangle).




Figure 1: ROI parameters for rectangular and elliptical geometries.

In the subband domain a ROI selection generates a corresponding binary subband mask $M(u, v)$ ((u, v) are the subband coefficient coordinates). The image-subband domain mask conversion is not univocal: in order to actually represent all the ROI pixels with an uniform (eventually lossless) reconstruction quality, without any degradation on the ROI contour, the action of the synthesis filters in the subsequent image reconstruction

is usually taken into account,^{1,5} but other solutions could also be adopted^{10,14} by exploiting the scale-space properties of the multiresolution representation. In the following the above conversion is performed by the binary operator SMG (subband mask generator) such that $M(u, v) = SMG(ROI(Ge, x, y))$ (or, in an abbreviated notation, $M(u, v) = SMG(Ge)$) is 0 for the subband BG and 1 for the subband ROI area, as shown in Fig.2. To indicate the relative importance of the ROI with respect to the BG a “saliency” factor Sa has to be defined, on a certain scale, to represent the user wish. Therefore the Sa parameter must be converted in a subband coefficient scaling factor s which determines the quality gap between the ROI and BG reconstruction quality. The scaling factor is usually chosen to be a power of two, $s = 2^N$, where N actually represents the number of bit-plane up-shifts for the ROI coefficients in the subband domain.⁵ This scaling policy turns to be computationally effective because wavelet coefficient are usually represented on integers, and does not represent a limitation for practical use. A scaling mask S can then be defined in the wavelet domain as $S(u, v) = N \cdot M(u, v)$. The presented ROI formalism has a corresponding syntax in JPEG2000 part 2, where multiple ROI can also be defined.

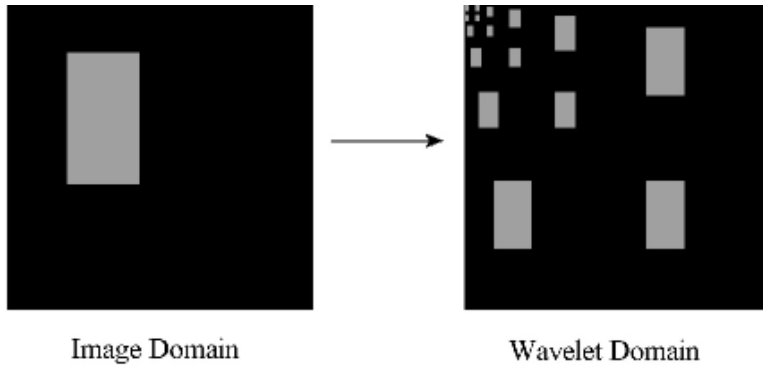


Figure 2: Binary masks: the image domain $ROI(Ge, x, y)$ and the subband domain $M(u, v)$.

Within this framework we present our solution to control the ROI quality decay. This control should correspond to perceptual or semantic issues and needs a second ROI qualifying parameter that we call “focalization”,¹¹ Fo . The Fo is expressed, on a certain scale, by the user or eventually by an automatic procedure. An high Fo value corresponds to a sharp quality transition need, on the contrary a low Fo signifies that the interest is not strongly localized into the ROI region, therefore the ROI neighborhoods must be gradually degraded. We propose to approximate the quality decay by a tiered pyramid of concentric ROI’s, with the original ROI on the top. In our bit-plane coefficient scaling framework the tier levels are a set of decreasing bit-plane shifts. The number of tiers is then implicitly defined by the number N and the pyramid should be constructed using a set of N ROI, which geometry vector is indicated with $Ge^{(n)}$, where $n \in [1, N]$, and the Fo is rendered, by means of a mask parameter f , as the size ratio between the top (N) and bottom (1) tiers of the pyramid: the highest focalization is represented by $f = 0$, while in general $f = \frac{a^{(1)} - a^{(N)}}{2a^{(N)}}$, where $a^{(N)} = a$ and $a^{(1)}$ is the smaller diameter of the greater ROI. This way, when $f = 1$ the pyramid lateral extension $\frac{a^{(1)} - a^{(N)}}{2}$ has the same size of the original ROI smaller diameter. Given the user defined triplet (Ge, Sa, Fo) , such that $N > 1$ and $f > 0$, a set of ROI is then defined for $n = 1, \dots, N$:

$$Ge^{(n)} : \begin{cases} x_t^{(n)} &= x_t - a \cdot f \cdot ((n-1)/(N-1)) \\ y_t^{(n)} &= y_t - a \cdot f \cdot ((n-1)/(N-1)) \\ x_b^{(n)} &= x_b + a \cdot f \cdot ((n-1)/(N-1)) \\ y_b^{(n)} &= y_b + a \cdot f \cdot ((n-1)/(N-1)) \\ t &= \text{rectangle, ellipse} \end{cases} \quad (1)$$

and the subband scaling mask is obtained by:

$$S(u, v) = \bigvee_{n=1}^N \left[n \cdot SMG(Ge^{(n)}) \right] \quad (2)$$

The \vee operator is the componentwise matrix *maximum* which generates the subband pyramids. Fig.3 shows different mask pyramids for some values of f while Fig.4 actually shows the action of eqs.(1) and (2) on the user defined (Ge, Sa, Fo) input. As an example, we selectively code the Lena image* using the “scaling-based”




Figure 3: Different ROI stacks in function of the f parameter.




Figure 4: Image domain $ROI(Ge^{(n)}, x, y)$ pyramid and subband scaling mask $S(u, v)$.

approach and the above multiple concentric ROI scaling, with a rectangular ROI selection. In Fig.5 the original image $I_{x,y}$ (a) is coded at 0.0625 bpp, with a conventional (b) and a concentric ROI (c) selective coding. In each case a scaling factor $s = 2^4$ ($N = 4$) has been used, while in (c) the quality decay is controlled by a factor $f = 0.5$. The images clearly shows the effects of the quality decay control. Fig.6 shows the $PSNR(x)$ evaluated, on this example, along the rows of the decoded images $\hat{I}_{x,y}$ and averaged on the columns where the rectangular ROI is present:

$$PSNR_{Ge}(x) = 10 \log_{10} \frac{255^2}{\frac{1}{y_b - y_t + 1} \sum_{y=y_t}^{y_b} (I_{x,y} - \hat{I}_{x,y})^2} \quad (3)$$

Even if the local information content obviously modulates the overall $PSNR_{Ge}(x)$ trend, the graph clearly shows the quality decay difference between the two ROI modalities. The graph also shows a predictable slight reduction in the ROI quality for the $f \neq 0$ case at a fixed bit-rate. This reduction however is not visible for high quality ROI reconstructions. On the other hand, if a lossless (or near lossless for non integer wavelet

*The presented techniques are not associated to a specific codec. In this paper, the EMDC progressive morphological wavelet coder presented in¹⁵ is used both for 2D and 3D coding.




Figure 5. Selective coding on Lena: (a) original image ROI selection, (b) sharp quality decay, (c) smoothed quality decay; (b) and (c) ROI boundary details.

transforms) ROI quality is exacted, the size of the coded bit-stream will be slightly superior in the case of concentric ROI, of a quantity that mainly depends from the size a and the factor f . In both cases, when the quality control gives an actual benefit in terms of image perception or interpretation, a slight difference in rate-distortion performance can be considered tolerable if not negligible. What has been done can be applied to a general wavelet coding technique and also perfectly reproduced with a JPEG2000 syntax, because multiple ROI can be defined in the standard where, in the case of ROI overlapping, the max-rule is applied.⁵ However,

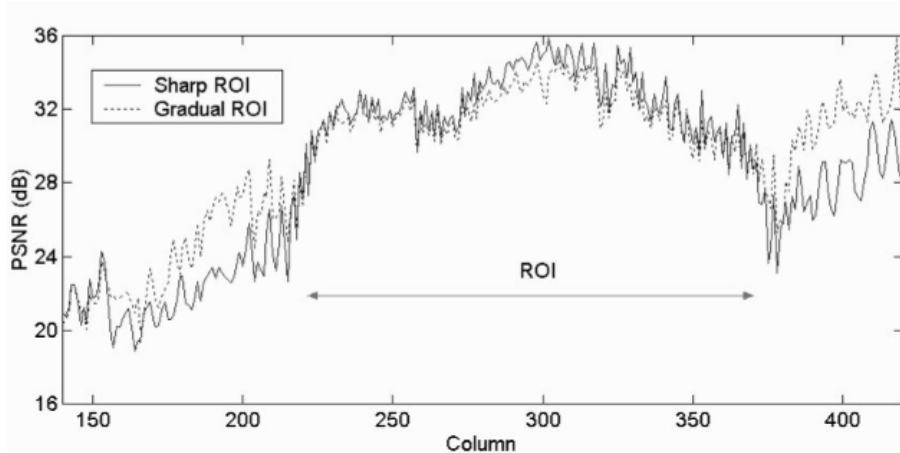


Figure 6: PSNR on the ROI raw (see eq.(3)) for the example of Fig.5.

being the quality decay management not explicitly envisaged by JPEG2000, a certain redundancy in the ROI pyramid description is necessary. In fact, even if only 7 parameters ($x_t, x_b, y_t, y_b, t, f, s$) would be sufficient to correctly build the $S(u, v)$ of eq.(2), following the JPEG2000 syntax (see ref⁵ annexes A and L) a set of $6 \cdot N$ parameters has to be declared ($5 \cdot N$ for $Ge^{(n)}$ and N scaling factors). In addition, JPEG2000 performs an optimized scaling for each coding block, but this is transparent to our technique.

3. VOLUME OF INTEREST DEFINITION AND DESCRIPTION

3D wavelet coding is directed to efficiently exploit the statistical properties of the 3D counterpart of the so called “natural image”,¹⁶ that we call “natural volume”. In other words, a still volume can be defined “natural” if it produces a “natural image” when sliced along an arbitrary direction. If this is not the case the efficiency of a wavelet 3D coding scheme drastically fall.¹⁷ Natural volumes are usually visualized slice by slice along one or more (usually orthogonal) slicing directions. In the biomedical field physicians are trained to make diagnosis on slices.

The need to code volumetric data is even more attracting, if not critical, due to the amount of data actually or potentially created by modern 3D scanners, especially in the biomedical field. In many practical situations, due to space restrictions in the archive, the acquired digital 3D datasets are subsampled (some representative slices are selected), printed on films and/or partially deleted after the diagnosis. Instead of this rough selective compression the volume of interest coding could be considered a better solution. However, a 3D extension of the selective coding reveals a set of problems in terms of definition and representation of volumes of interest (VOI). There is the issue of how effectively select a VOI on a volumetric dataset and, for the “scaling-based” approach, the problem to represent it in a flexible and compact parametric form. Moreover, the JP3D working draft¹² is oriented to “simply” extend of the already standardized 2D techniques without adding, if possible, too much different solutions. Selecting a VOI is not trivial and risks to be considered a problem related to the existence of advanced 3D user interfaces. However, the most part of 3D coding application scenarios would not require a precise segmentation of a 3D object prior to the coding (which is also impractical for everyday clinical use), but an effective and sufficiently flexible indication of the volume which contains the structures of diagnostic interest. From the user (physician) point of view, the VOI selection should be straight, simple and well-matched with its own culture and working (diagnostic) practice.

We propose a solution that satisfies all the above requirements and is based on the specific attributes of the volumetric data to be coded. We generate a VOI on the base of a set composition (intersection or union) of solids which are generated by rectangular or elliptical ROI selections on 2D volume slices, as provided by JPEG2000 part 2. With a simple graphical interface it is possible to effectively select a VOI by cutting the volume along the three orthogonal direction and selecting the opportune ROI on representative slices. This is shown in Fig.7,

where representative slices have been quickly found and the ROI shapes selected by conventional interaction tools. The three-dimensional subband VOI mask $S_{3D}(u, v, w)$ can be easily calculated by the following equation




Figure 7: VOI selection by a triple ROI definition.

$$S_{3D}(u, v, w) = N \cdot SMG_{3D} [COI_x(Ge_{yz}) \cap COI_y(Ge_{xz}) \cap COI_z(Ge_{xy})] \quad (4)$$

which realizes a cylindroid intersection. Unions or mixed compositions can be also envisaged. In the above equation SMG_{3D} is a simple 3D extension of the subband mask generator, while the cylindroids of interest (e.g. COI_x) are generated by prolonging each planar ROI (e.g. described by Ge_{yz}) along the corresponding orthogonal direction. Fig.8 shows an example of three intersecting COI and the result of their composition. With this technique the VOI is compactly represented by the parameters of 3 ROI's and the type of solid composition (intersection, union or mixed). The above technique can be perfectly combined with the progressive quality




Figure 8: VOI definition by cylindroids composition: original cylindroids (a) and their intersection (b).

decay described above. This could be useful for biomedical applications as shown in the following section. In this case the $S_{3D}(u, v, w)$ is generated by

$$S_{3D}(u, v, w) = \bigvee_{n=1}^N \left\{ n \cdot SMG_{3D} \left[COI_x(Ge_{yz}^{(n)}) \cap COI_y(Ge_{xz}^{(n)}) \cap COI_z(Ge_{xy}^{(n)}) \right] \right\} \quad (5)$$

Fig.9 shows an example of VOI generated with 3 rectangular ROI with scaling factor 2 and $f \neq 0$.

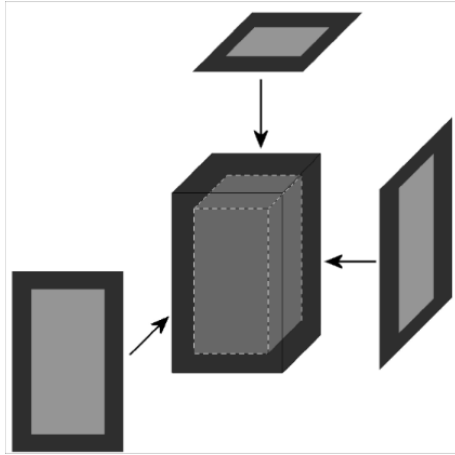


Figure 9: VOI with controlled quality decay.

4. DIAGNOSTIC IMAGING APPLICATION EXPERIMENTS

As mentioned above, the quality managing for medical imaging is a critical point. Diagnostic quality can be defined as the possibility to make the same diagnosis on both the original and the decoded images, and measured on a subjective scale.¹⁰ The efficacy of the concentric ROI approach in terms of improved diagnostic quality of ROI-decompressed images is here exemplified by means of ad-hoc experiments in collaboration with an hospital staff (see the acknowledgments). We analyze two cases of brain lesions on MR (Magnetic Resonance) images as representative of a wide range of diagnostic situations. Usually, brain lesions are localized masses that can be indicated and encircled by a neuroradiologist and selectively ROI coded with an adequate overall bit-rate. A smoothed ROI-BG quality decay guarantees that the neighborhood of the lesions are preserved from objectionable artifacts as well. In fact in those cases, even if the diagnostic interest is concentrated in the ROI area, the evaluation of the interactions between the lesion and the brain and the visual analysis of the surrounding area (interested for example by an edema or an hematoma) are essential to a correct diagnosis achievement. The results presented here confirm the above statements and should be taken as a starting point to justify and to plan a more extended clinical investigation. Fig.10 shows the selection of a ROI corresponding to a lesion (in this case an hematoma), and two enlarged views of the 0.0625 bpp decoded area of interest, with and without applying a conventional scaling-based selective coding with $N = 4$. In Fig.11 the conventional ROI coding is compared to the proposed quality decay management approach, where two value of f has been used. From a general (non-professional) point of view the observation made for the Lena example of Sec.2.2 still hold. In addition, the neuroradiologist judged the three decoded images: at the same overall bit-rate of 0.16 bpp the two images with $f \neq 0$ has been preferred (with a slight bias toward the higher f value) with respect to the conventional ROI decoded one. In fact, the latter has been classified with a fair diagnostic quality, while the former gained good diagnostic quality judgments[†].

In the second example the 3D selective coding is examined. The VOI selection of Fig.7 has been combined with the concentric ROI approach to perform a 3D quality decay management. A scaling factor $N = 4$ and different focalization parameters $f = 0$ (conventional), $f = 0.5$ and $f = 2$ have been used with a 0.16 bpp coding rate. In Fig.12 2D slices and 3D views of the decoded dataset are shown, where the 3D views have been obtained by direct volume rendering, using two cutting plane to localize the lesion. Even in this case the conventional and proposed approaches can be compared both visually and semantically. The professional judgment still was for the smoothed quality decay images.

[†]In Fig.11 an augmented compression ratio has been used to better render the visual result on printed paper.




Figure 10. ROI selection on a MR image; 0.0625 bpp reconstructed image details with ($N = 4$, $f = 0$) and without ROI coding.

5. CONCLUSIONS

The success of selective coding is mainly connected with a quick and flexible ROI management. In some professional environment, the relations between selective coding potential benefits and working environment constraints and practice are of crucial importance. In this paper some new selective coding ideas have been presented: a) a smoothed ROI-BG quality decay control has been introduced in a multiple ROI framework; b) a 3D extension of the ROI coding has been proposed for 3D datasets. The two ideas can be combined and set in a JPEG2000 environment. In particular, the quality control can be implemented using a JPEG2000 syntax, while the 3D extension could be considered for the JP3D ongoing standardization. From a user point of view the ROI is described and qualified by the vector (Ge, Sa, Fo) or by $(Ge_{yz}, Ge_{xz}, Ge_{xy}, Sa, Fo)$ in the case of a 3D VOI. The focalization Fo should express a general user need (perceptual, aesthetic,...) or a more precise exigency to preserve the semantic meaning of the images, as in diagnostic imaging. The 3D VOI definition is simple and seems to be well matched to the conventional ROI selection and visualization modalities for “natural” (not too anisotropic) volumetric datasets, therefore it is well suited to be implemented with good usability at a human-machine interface level: this is a general and essential pre-condition for the acceptance of an image processing tool in a clinical environment.

ACKNOWLEDGMENTS

Prof. Roberto Gasparotti, head of the Neuroradiology division of the Hospital “Spedali Civili di Brescia” and Associate Professor at the University of Brescia, for his fruitful contribution to this work.

REFERENCES

1. D. Nister and C. Christopoulos, “Lossless region of interest coding,” *Signal Processing* **78**, pp. 1–17, 1999.
2. K. h. Park and H. W. Park, “Region-of-interest coding based on set partitioning in hierarchical trees,” *IEEE Trans. on Circ. and Syst. for Video Technol.* **12**, pp. 106–113, Feb. 2002.
3. J. Askelöf, M. L. Carlander, and C. Christopoulos, “Region of interest coding in JPEG 2000,” *Signal Proc.: Image Comm.* **17**, pp. 105–111, 2002.
4. ISO/IEC 15444-1, “Information technology - JPEG2000 image coding system - Part 1: Core coding system.” Final Publication Draft, 2002.




Figure 11. ROI coding (0.0625bpp) with $N = 4$ and different f values. Decoded MR images and corresponding ROI details: (a) $f = 0$, (b) $f = 0.5$, (c) $f = 1.0$.

5. ISO/IEC 15444-2, "Information technology - JPEG2000 image coding system - Part 2: Extensions." Final Publication Draft, 2002.
6. ISO/IEC WD15444-9, "Information technology - JPEG2000 image coding system - Part 9: Interactivity tools, APIs and protocols." Working Draft, 2002.
7. B. J. Erickson, "Irreversible compression of medical images." white paper from SCAR (Society for Computer Applications in Radiology), www.scarnet.org, Nov. 2000.
8. J. Ström and P. Cosman, "Medical image compression with lossless regions of interest," *Signal Processing* **59**, pp. 155–171, June 1997.
9. A. Signoroni and R. Leonardi, "Progressive ROI coding and diagnostic quality for medical image compression," in *VCIP'98*, **3309**, pp. 674–685, SPIE, (San José, CA), Jan. 1998.




Figure 12. Detail images from a 0.16 bpp decoded volume with VOI ($N = 4$). 2D sections (slices): (a) $f = 0$, (b) $f = 0.5$ and (c) $f = 2$. Direct volume rendering (using two cutting planes): (d) $f = 0$, (e) $f = 2$.

10. A. Signoroni and R. Leonardi, "Progressive medical image compression using a diagnostic quality measure on regions-of-interest," in *EUSIPCO'98*, pp. 2325–2328, (RODI,Gr), Sept. 1998.
11. A. Signoroni, A. Bianchi, P. Mussio, and R. Leonardi, "Design of a visual environment for evaluating and customizing medical image compression techniques," in *ICIAP'99*, pp. 1021–1026, (Venice,I), Sept. 1999.
12. ISO/IEC WD15444-10, "Information technology - JPEG2000 image coding system - Part 10: Extensions for three-dimensional data and floating point data." Working Draft, 2002.
13. E. Nguyen Phuc, *Compression sélective et focalisation visuelle: application au codage hybride de séquences d'images*. PhD thesis, Université de Rennes 1, F, 1995.
14. Z. Wang and A. C. Bovik, "Embedded foveation image coding," *IEEE Trans. on Image Processing* **10**, pp. 1397–1410, Oct. 2001.
15. F. Lazzaroni, A. Signoroni, and R. Leonardi, "Embedded morphological dilation coding for 2D and 3D images," in *VCIP 2002*, **4671**, pp. 923–934, SPIE, (San José, CA), Jan. 2002.
16. E. P. Simoncelli and B. A. Olshausen, "Natural image statistics and neural representation," *Annu. Rev. Neurosci.* **24**, pp. 1193–1216, May 2001.
17. A. Signoroni and R. Leonardi, "Diagnostic compression of biomedical volumes," in *EUSIPCO 2000*, **I**, pp. 533–536, (Tampere, Finland), Sep. 2000.