

Segmentation-based lossless compression of burn wound images

Carmen Serrano

Begoña Acha Piñero

Área de Teoría de la Señal y Comunicaciones

Escuela Superior de Ingenieros

Universidad de Sevilla

Camino de los Descubrimientos

s/n. 41092 Sevilla, Spain

E-mail: cserrano@viento.us.es, bacha@viento.us.es

Rangaraj M. Rangayyan

Department of Electrical and Computer Engineering

University of Calgary

Calgary Alberta, T2N 1N4, Canada

Laura M. Roa

Grupo de Ingeniería Biomédica

Escuela Superior de Ingenieros

Universidad de Sevilla

Camino de los Descubrimientos

s/n. 41092 Sevilla Spain

Abstract. *Color images may be encoded by using a gray-scale image compression technique on each of the three color planes. Such an approach, however, does not take advantage of the correlation existing between the color planes. In this paper, a new segmentation-based lossless compression method is proposed for color images. The method exploits the correlation existing among the three color planes by treating each pixel as a vector of three components, performing region growing and difference operations using the vectors, and applying a color coordinate transformation. The method performed better than the Joint Photographic Experts Group (JPEG) standard by an average of 3.40 bits/pixel with a database including four natural color images of scenery, four images of burn wounds, and four fractal images, and it outperformed the Joint Bi-Level Image experts Group (JBIG) standard by an average of 3.01 bits/pixel. When applied to a database of 20 burn wound images, the 24 bits/pixel images were efficiently compressed to 4.79 bits/pixel, then requiring 4.16 bits/pixel less than JPEG and 5.41 bits/pixel less than JBIG. © 2001 SPIE and IS&T.*

[DOI: 10.1117/1.1383781]

1 Introduction

Lossless compression techniques are essential in archival and communication of medical images. Increasing applications of telemedicine in the areas of dermatology and pathology are creating new demands, such as transmission and archival of color images. The Burn Unit of the Hospital Universitario Virgen del Rocío de Sevilla and the Biomedical Engineering and Signal Processing Groups of the University of Sevilla are developing a telemedicine project

where the diagnosis of burn patients is performed with digital color photographs. As automatic classification is performed in a procedure where color and texture are essential, a lossless compression method is needed for this particular application.

Although significant effort has been directed towards the development of lossless algorithms for compressing image data, most of such methods have been oriented towards compressing gray-scale or two-tone (binary) images. It is commonly stated that a red-green-blue (RGB) color image can be easily compressed by using a gray-scale image compression technique on each of the three color components. However, such an approach does not take advantage of the correlation existing between the color planes.

Recently, a few efforts have been directed towards compression of true-color images by taking advantage of their spectral correlation Singh *et al.*¹ proposed an integer-based wavelet transform method, which was shown to be efficient for compression of images with fine details. Memon and Sayood² proposed several methods based on error prediction models, and reported improvement over the performance of the Joint Photographic Experts Group (JPEG) standard of about 1.5 bits per color pixel. Bockstein³ proposed a new method based on a lossless transformation from the RGB planes to other planes, defined as linear combinations of the former, and compressing each of the new planes with a gray-scale image lossless compression algorithm.⁴ Van Assche *et al.*⁵ proposed a technique based on the Karhunen–Loève transform (KLT) to decorrelate the color planes; compression rates of about 0.5–2.0 bits per pixel better than those provided by lossless JPEG were ob-

Paper 99083 received Dec. 22, 1999; revised manuscript received Aug. 7, 2000; accepted for publication Feb. 23, 2001.
1017-9909/2001/\$15.00 © 2001 SPIE and IS&T.

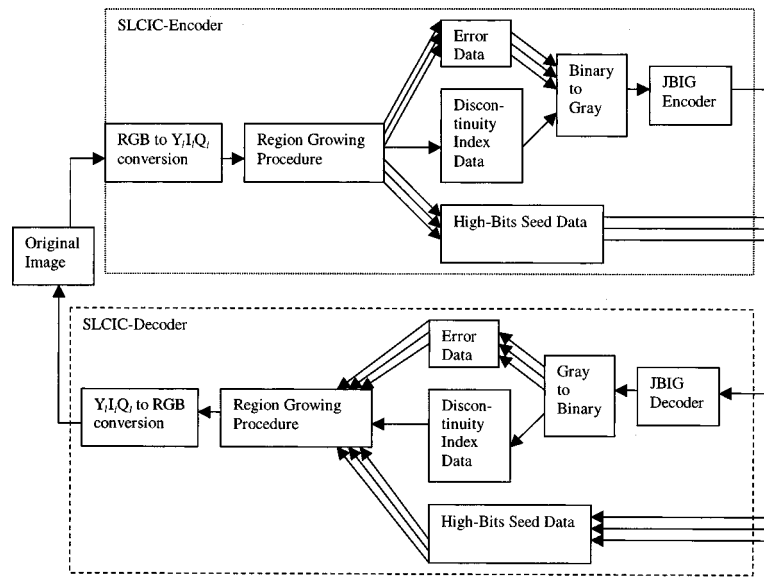


Fig. 1 Summary of the SLCIC procedure.

tained, but with the disadvantage of high computational time.

In 1985, Kunt *et al.*⁶ proposed a contour-texture approach to picture coding; they called such approaches second-generation image coding techniques, and they applied it to lossy compression. The main idea behind this technique is to first segment the image into nearly homogeneous regions surrounded by contours such that the contours correspond, as much as possible, to those of the objects in the image, and then to encode the contour and texture information separately. Because contours can be represented as one-dimensional signals and pixels within a region are highly correlated, such methods are expected to attain high compression ratios. Although the idea seems to be promising, its implementation meets a series of difficulties. A major problem exists at its very important first step of segmentation, which determines the final performance of the segmentation-based coding method. Most of the segmentation algorithms based upon this principle are sophisticated and give good performance only for specific types of images.

Shen and Rangayyan⁷ proposed a segmentation-based lossless image coding (SLIC) technique for radiographic images. To overcome the problem mentioned above in relation to segmentation, they proposed a simple region growing method. Instead of generating a contour set, a discontinuity map is obtained during the region growing procedure. Concurrently, the method also provides a corresponding error image based on the difference between each pixel and its corresponding "center pixel." The last step in the algorithm is to code the discontinuity and error data with the Joint Bilevel Image experts Group (JBIG) compression standard.⁸

We have recently extended the SLIC algorithm to color images.⁹ In order to exploit the correlation existing among the three color planes, our algorithm treats each pixel of the image as a three-dimensional (3D) vector (RGB) and performs 3D region growing. The method produces a three-component error image but only a one-component discontinuity map.

With a view to obtain better compression, we have included a lossless color-coordinate conversion step from RGB to the luminance, in-phase, and quadrature-phase (YIQ) representation system.¹⁰

We devote Sec. 2 to provide details of the segmentation-based lossless color image coding (SLCIC) algorithm, and provide the results of its comparison with the JPEG international standard for lossless still-image compression¹¹ in Sec. 3.

2 Segmentation-Based Lossless Coding Algorithm for Color Images

In lossless image compression, the task is usually split into two stages: one is image transformation, with the purpose of data decorrelation; the other is encoding of the transformed data. In the SLCIC algorithm, image transformation is achieved in both the region growing procedure and, later, in the JBIG algorithm. The JBIG algorithm uses an adaptive three-dimensional coding model followed by an adaptive arithmetic coder.⁸ Figure 1 gives a block diagram demonstrating the various steps in the SLCIC procedure.

2.1 Region Growing Algorithm

As described by Shen and Rangayyan⁷ for single-component images, the region growing procedure starts with a single pixel, called the seed pixel. Each of the seed's four-connected neighbor pixels, scanned in a fixed order, is checked with a region growing (or inclusion) condition. If the condition is satisfied, the neighbor pixel is included in the region. The procedure is recursively continued until no spatially connected pixel meets the growing condition. A new region growing procedure is then started with the next pixel of the image which is not already a member of a region. The procedure ends when every pixel in the image has been included in one of the regions grown.

The region growing conditions used in this work are (i) the neighbor pixel is not a member of any of the regions already grown, and (ii) the absolute difference between the

red, green, and blue intensities of the neighbor pixel and those of the corresponding “center pixel” is less than a predefined $error_level$. The center pixel is defined as that pixel which is being used as the reference to check its four-connected neighbors for inclusion in the region being grown; the center pixel would have already been included in the region. When a new neighbor pixel is included in the region being grown, its $error_level$ -shift-up difference with respect to its center pixel is stored as the pixel’s “error” value; this value is a vector of three components. If only the first of the two region growing conditions is met, the discontinuity index of the pixel is incremented. The discontinuity index indicates how many times a pixel has been tested to be included in a region and finally it has not been included. Therefore, after region growing, a “discontinuity index image data part” and a three-component “error image data part” will be obtained. The maximum value of the discontinuity index is four. Most of the previously reported segmentation-based coding algorithms include contour coding and region coding; instead of these steps, we use a discontinuity index data part and an error image data part.

The $error_level$ is determined by the preselected $error_bits$ assigned to each of the three components of the vector as $error_level = 2^{error_bits - 1}$. For instance, if the error image is allowed to take up to 5 b/color-plane-pixel for each component ($error_bits = 5$), the corresponding $error_level$ is 16; in other words, the allowed difference range is then $[-15, 15]$. The error value of the seed pixel of each region is defined as the vector formed by the low $error_bits$ bits of each component. The value of the high $(N - error_bits)$ bits of the three components of the pixel is stored in a “high-bits seed data part,” where N is the number of b/pixel in the original image data.

The above three data parts are used to fully recover the original image during the decoding process. The region growing conditions during decoding are (i) the neighbor pixel under consideration for inclusion is not in any of the previously grown regions, and (ii) the discontinuity index of the pixel is equal to zero. When the conditions are met for a pixel, its pixel value is restored as the sum of its $error_level$ -shift-down error value and its center pixel value (except for the seed pixel of every region, for which the latter is given by the high-bits seed data part). If only the first of the two conditions is satisfied, the discontinuity index of the pixel is decremented. Thus the discontinuity index generated during segmentation is used to guide region growing during decoding.

The region growing procedure may be viewed as adaptive scanning of the image with the aim of maintaining a localized difference or error value within a limited dynamic range prespecified and controlled by $error_bits$.

Figure 2 is a simple example for illustration of the region growing procedure and its result. The original image is the 512×512 24-bit Lena image. A 5×5 segment of the image is shown in Fig. 2(a). The value of $error_bits$ is set to three in this example. Figure 2(b) is the result of region growing. The corresponding three data parts, namely discontinuity index image data, error image data, and high-bits seed data, are shown in Figs. 2(c)–2(e).

46 114 151	45 115 151	55 110 151	84 176 184	55 111 153
45 115 151	45 115 151	55 110 151	52 112 152	55 111 153
45 115 154	45 115 154	51 110 151	48 110 154	55 110 151
48 114 152	48 114 152	50 110 151	52 115 153	54 113 152
55 114 151	55 114 151	49 116 152	48 111 152	55 112 152

(a)

46 114 151 (seed)	45 115 151	55 110 151 (seed)	84 176 184 (seed)	55 111 153
45 115 151	45 115 151	55 110 151	52 112 152	55 111 153
45 115 154	45 115 154	51 110 151 (seed)	48 110 154	55 110 151
48 114 152	48 114 152	50 110 151	52 115 153	54 113 152
55 114 151 (seed)	55 114 151	49 116 152 (seed)	48 111 152 (seed)	55 112 152

(b)

0 0 1 3 0				
0 0 1 0 0				
0 0 2 3 0				
0 0 2 0 0				
1 1 2 3 0				

(c)

6 2 7	3 5 4	7 6 7	4 0 0	4 4 4
3 5 4	4 4 4	4 4 4	1 6 5	7 3 5
4 4 7	4 4 7	3 6 7	1 4 7	4 3 2
7 3 2	7 3 2	3 4 4	2 6 5	3 7 5
7 2 7	4 4 4	1 4 0	0 7 0	5 3 4

(d)

5 14 17	6 13 17	10 22 23	6 13 17	6 14 17	6 14 35	6 13 35
---------	---------	----------	---------	---------	---------	---------

(e)

Fig. 2 Simple example of the region growing procedure and its result with $error_bits$ set to be three: (a) a 5×5 segment of the 512×512 24-bit color Lena image, where each value represents the R, G, and B component, respectively; (b) result of region growing; (c) corresponding discontinuity index data part; (d) corresponding error image data part; (e) the corresponding high-bits seed data part.

2.2 Segmentation-Based Lossless Color Image Compression Procedure

The complete SLCIC procedure is illustrated in Fig. 1. At the encoding end, the original image is transformed by the region growing procedure into three parts: the discontinuity index image data part with one-dimensional elements, and the error image data part and the high-bits seed data part composed of 3D vectors. The first two data parts are then gray coded, broken down into bit planes, and finally JBIG coded. The last data part is stored or transmitted as-is; it needs $(N - error_bits) * 3$ bits per seed pixel.

At the receiving end JBIG-coded data files are JBIG-decoded first and then the gray-coded bit planes are composed back to binary code. Finally, the parts are combined together by the same region growing procedure to recover as in the encoding scheme to recover the original image.

2.3 Lossless RGB-to-YIQ Transformation

Usually, when coding color images with loss, a transformation from RGB color coordinates to YIQ is performed. This is based upon the observation that the Y component is uncorrelated with the chrominance components (I and Q), and that most high-frequency components of a color image are concentrated in the Y component.¹² In lossless compression, the second advantage cannot be exploited, but the first one is used to gain improved compression.

To change the color coordinates from RGB to YIQ the following linear transformation is performed¹²:

$$\begin{bmatrix} Y \\ I \\ Q \end{bmatrix} = \begin{bmatrix} 0.299 & 0.587 & 0.114 \\ 0.596 & -0.274 & -0.322 \\ 0.211 & -0.523 & 0.312 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}. \quad (1)$$

However, this transformation cannot be applied in lossless compression because it involves floating-point operations, thus requiring quantization. In order to avoid quantization error, a transformation based on integer arithmetic is required. Such a transformation is completely reversible and hence lossless, and is given by:¹

$$Y_l = \left\lfloor \frac{\left\lfloor \frac{R+B}{2} \right\rfloor + G}{2} \right\rfloor, \quad (2)$$

$$I_l = R - B, \quad (3)$$

$$Q_l = \left\lfloor \frac{R+B}{2} \right\rfloor - G, \quad (4)$$

where $\lfloor \cdot \rfloor$ denotes the floor function.

The inverse transformation is given by:

$$R = Y_l + \left\lfloor \frac{Q_l + 1}{2} \right\rfloor + \left\lfloor \frac{I_l + 1}{2} \right\rfloor, \quad (5)$$

$$G = Y_l - \left\lfloor \frac{Q_l}{2} \right\rfloor, \quad (6)$$

$$B = Y_l + \left\lfloor \frac{Q_l + 1}{2} \right\rfloor - \left\lfloor \frac{I_l}{2} \right\rfloor. \quad (7)$$

The subscript l is used to indicate the lossless nature of the transformation. This transformation performs the YIQ one, but without losses, and it is proved in Ref. 1. It is important to note that in order to preserve the lossless nature of the algorithm, one bit must be added to each of the chrominance error components (one for I and one for Q) due to the sign that appears when converting from RGB to $Y_l I_l Q_l$.

3 Experimental Results

The SLCIC technique was tested using 12 24-bit color images of different sizes and nature. The performance of the SLCIC technique was compared with that of the lossless JPEG 7 standard and the JBIG standard. The initial test image set includes four images of scenery (Lena, Peppers, Parrots, and Germany), four images of burn wounds, and four fractal images (from http://dir.yahoo.com/Arts/Visual_Arts/Computer_Generated/Fractals/Artists/). The burn wound images have been taken following a protocol explained in Ref. 13. In this protocol a digital photograph camera is used, and the photographs are stored as true-color ones, i.e., 24 bits per pixel, 8 bits per each color component (RGB: red, green, blue).

Figures 3–5 show the error and the discontinuity index data parts as images for three of the test images used. It can be observed in all of the cases illustrated that the error image does not contain any significant chromaticity infor-

mation (actually they look like gray-scale images being color ones), but retains a portion of the boundary information of the image. The discontinuity index clearly corresponds to edges in the image.

The compression performance of the SLCIC method with the initial set of 12 images is summarized in Tables 1 and 2 along with the results of JPEG and JBIG compression. The SLCIC method has a tunable parameter, which is *error_bits*. The SLCIC method has outperformed lossless JPEG on the average for the test image set used by 0.68 bits per color pixel (with the original images having 24 b/pixel in the RGB domain), and JBIG by 0.29 b/pixel. As shown in Table 1, when using the $Y_l I_l Q_l$ coordinates, the performance of the SLCIC is much better, providing an average of 9.24 b/pixel. The method outperformed standard JPEG by an average of 3.34 b/pixel, JBIG by 2.95 b/pixel, and SLCIC with RGB coordinates by 2.66 b/pixel. In Table 2 a comparison of the code and decode times for SLCIC, JPEG, and JBIG is done, where the times of reading and writing the images to and from the disk are also considered and, for the case of the SLCIC algorithm, the time employed in the conversion from RGB to $Y_l I_l Q_l$ is also included. Although internally SLCIC uses JBIG, it gets faster times than JBIG alone. This is due to the fact that JBIG acting alone has to code 24 bit planes whereas JBIG joint to SLCIC only has to code the number of bit planes determined by the parameter *error_bits*. Additionally, the region growing procedure does not add a significant quantity to the total computational time. The JBIG software we have

Table 1 Comparison of the performances of SLCIC, JBIG, and JPEG using twelve 24 b images by b/pixel.

Image	JPEG (b/pixel)	JBIG (b/pixel)	SLCIC (b/pix)	
			RGB	$Y_l I_l Q_l$
Scenery:				
1. Peppers	15.69	15.65	15.99	16.29
2. Germany	17.73	17.49	16.49	12.35
3. Lena	14.79	15.12	15.08	14.98
4. Parrots	18.04	12.06	18.01	9.63
Average for scenery:	16.56	15.08	16.39	13.31
Burn wounds:				
1. Photo 1	10.8	12.43	10.64	5.87
2. Photo 2	7.36	8.19	6.56	4.04
3. Photo 3	7.45	8.68	6.8	4.32
4. Photo 4	7.33	8.67	6.74	4.22
Average for burn wounds:	8.24	9.49	7.69	4.61
Fractals:				
1. Atlantis	15.12	14.05	13.76	12.76
2. 7 thm	12.09	11.27	11.00	10.97
3. Mandilv	14.58	14.46	13.45	7.90
4. Yinywthm	9.95	8.17	8.25	7.62
Average for fractals:	12.94	11.99	11.62	9.81
Average for all images:	12.58	12.19	11.9	9.24




Fig. 3 Lena: (a) original image; (b) error image; (c) discontinuity index image.




Fig. 4 Burn wound: (a) original image; (b) error image; (c) discontinuity index image.




Fig. 5 Fractal: (a) original image; (b) error image; (c) discontinuity index image.

Table 2 Comparison of code and decode time for JPEG, JBIG, and SLCIC-Y,I,Q. Computing time of less than 1 s is entered as 0 in the table.

Image	Size (col×row×b/pixel)	Code/decode time (s)		
		JPEG	JBIG	SLCIC
Scenery:				
1. Peppers	512×512×24	2/0	14/14	12/12
2. Germany	768×512×24	2/0	19/19	18/18
3. Lena	512×512×24	2/0	13/13	12/12
4. Parrots	768×512×24	3/0	19/19	16/16
Burn wounds:				
1. Photo 1	832×624×24	4/1	23/23	17/17
2. Photo 2	832×624×24	4/1	22/22	19/19
3. Photo 3	832×624×24	5/1	22/22	17/17
4. Photo 4	832×624×24	5/1	21/21	18/18
Fractals:				
1. Atlantis	320×240×24	0/0	4/4	1/1
2. 7 thm	160×120×24	0/0	1/1	0/0
3. Mandsilv	160×159×24	0/0	1/1	0/0
4. Yinywthm	320×240×24	0/0	4/4	1/1

employed is a free distribution version developed by AT&T in 1991.

After examining the Y, I, and Q component images separately, we realized that the I and Q components were much more uniform than the luminance one so, they would need less *error_bits* to be encoded than the Y component. We changed the algorithm in order to code each component with different number of *error_bits*, and we got better results, that are summarized in Table 3. To avoid confusions we have called this modification of the SLCIC algorithm, SLCIC2. With the SLCIC2 the improvement gained over JPEG is 3.40 b/pixel on average and 3.01 b/pixel over JBIG.

Due to the good results obtained with burn wound images, we were encouraged to study the algorithm for our particular application. Hence, we tested the SLCIC procedure with a larger database of 20 burn wound images. The results are summarized in Table 4, given as the mean and the variance for SLCIC, SLCIC2, JPEG, and JBIG, where we can see that the SLCIC method has outperformed standard JPEG by an average of 4.13 b/pixel and JBIG by 5.38 b/pixel, and the SLCIC2 has outperformed JPEG by an average of 4.16 b/pixel and JBIG by 5.41 b/pixel. The reason for such a good performance with burn wound images can be found in the special features of the images: they have uniform colors, such that majority of its information is located in the luminance component exclusively. Additionally, even in the luminance component, there are large, almost-uniform regions.

In the same way as the error matrix and the discontinuity index matrix are coded with JBIG after the region growing step, we experimented with encoding the seed vector with

Table 3 Comparison of the performances of SLCIC2, JBIG, and JPEG using twelve 24 b images by b/pixel. The SLCIC2 algorithm is implemented using different number of *error_bits* for the Y, I, and Q components.

Image	JPEG (b/pixel)	JBIG (b/pixel)	SLCIC2 (b/pix)
Scenery:			
1. Peppers	15.69	15.65	16.29
2. Germany	17.73	17.49	12.32
3. Lena	14.79	15.12	14.97
4. Parrots	18.04	12.06	9.61
Average for scenery:	16.56	15.08	13.30
Burn wounds:			
1. Photo 1	10.8	12.43	5.83
2. Photo 2	7.36	8.19	4.01
3. Photo 3	7.45	8.68	4.29
4. Photo 4	7.33	8.67	4.20
Average for burn wounds:	8.24	9.49	4.58
Fractals:			
1. Atlantis	15.12	14.05	12.69
2. 7 thm	12.09	11.27	10.80
3. Mandsilv	14.58	14.46	7.58
4. Yinywthm	9.95	8.17	7.55
Average for fractals:	12.94	11.99	9.66
Average for all images:	12.58	12.19	9.18

JBIG. However, the results were poor because the number of regions obtained was not large.

4 Discussion

While there have been some works on lossless compression of color images, very few have exploited the notion of what has been called second-generation image coding,¹⁴ in particular the segmentation-based approach. In this article, a method based on segmentation (region growing) has been presented. When tested with several general-purpose images, the proposed method resulted in an improvement of 3.40 bits per color pixel over the lossless JPEG standard and of 3.01 bits per color pixel over the JBIG standard. When tested with a database of 20 burn wound images, the method outperformed standard JPEG by 4.16 bits per color pixel, and standard JBIG by 5.41 bits per color pixel. As we said, the reason for such a good performance with burn wound images can be found in the special features of the images: they have uniform colors, such that majority of its information is located in the luminance component exclusively. Additionally, even in the luminance component, there are large, almost-uniform regions. So, we get very good results when considering different number of *error_bits* for each component. The chrominance components (I and Q) need less *error_bits* than the luminance one, because they carry less information.

The computational speed of SLCIC is not very slow comparing with that of JPEG due to the use of an efficient

Table 4 Comparison of the performances of SLCIC, SLCIC2, JPEG, and JBIG by b/pixel using 20 burn wound images. Each image is of size 832×624 pixels, with 24 b/pixel in the RGB representation.

20 burn wound images	JPEG (b/pixel)	JBIG (b/pixel)	SLCIC (b/pixel)	SLCIC2 (b/pixel)
Average	8.95	10.20	4.82	4.79
Variance	2.16	2.07	0.54	0.54

region growing method. Most of the computational cost of SLCIC is due to JBIG coding included in the algorithm. The region growing procedure does not add a significant quantity to the total computational time. In a SUN Ultra SPARC (120 MB memory, 2.1 GB disk, clock 167 MHz, 112 KB cache memory) SLCIC required 1–20 s to encode/decode the images used; JPEG required 1–5 s for the same images.

We are exploring possibilities of further improvement in the performance of SLCIC by analyzing the nature of the error image. One improvement under consideration is, instead of changing color coordinates of the whole image, to change the representation of only the error image from the RGB planes to hue, saturation, and value (HSV) or to $Y_I I_Q$. We are also exploring the possibility of using techniques other than JBIG, after the region growing step, to code the error and discontinuity parts in order to take into account the individual characteristics of each part.

Acknowledgments

This work was supported by grants from CICYT (TIC-96-0500-C10-08) and Junta de Andalucía of Spain and the Natural Sciences and Engineering Research Council (NSERC) of Canada. We thank Liang Shen for help with the segmentation-based lossless image coding technique for monochrome images, and Mihai Ciuc for assistance with color image processing. We also thank Dr. Tomás Gómez-Cía of the Unidad de Quemados of the Hospital Universitario Virgen del Rocío de Sevilla for providing the burn wound images.

References

1. I. Singh, P. Agathoklis, and A. Antoniou, "Lossless compression of color images using an improved integer-based nonlinear wavelet transform," *Proc. IEEE Int. Symp. Circuits Syst.*, 2609–2612 (1997).
2. N. D. Memon and K. Sayood, "Lossless compression of RGB color images," *Opt. Eng.* **34**(6), 1711–1717 (1995).
3. I. M. Bockstein, "An adaptive method of lossless compression of color images," *Proc. SPIE. Image Proc. Comput. Opt.*, **2363**, 300–306 (1994).
4. I. M. Bockstein, "A method of lossless image compression," in *Pattern Recognition and Image Analysis*, W. Poelzleitner and E. Wenger, Eds., pp. 92–98, Schriftenreihe der Oesterreichischen Computer Gesellschaft, Wien, Austria (1993).
5. S. Van Assche, W. Philips, and I. Lemahieu, "Lossless compression of pre-press images using linear color decorrelation," *Proc. IEEE Data Compression Conf.*, 578 (1998).
6. M. Kunt, A. Ikononopoulos, and M. Kocher, "Second-generation image-coding techniques," *Proc. IEEE* **73**(4), 549–574 (1985).
7. L. Shen, and R. M. Rangayyan, "A segmentation-based lossless image coding method for high-resolution medical image compression," *IEEE Trans. Med. Imaging* **16**(3), 301–307 (1997).
8. "Progressive bi-level image compression," International Telegraph and Telephone Consultative Committee (CCITT). Recommendation T. 82, 1993.


9. B. Acha, C. Serrano, R. M. Rangayyan, and L. M. Roa, "Lossless compression algorithm for colour images," *Electron. Lett.* **35**(3), 214–215 (1999).
10. C. Serrano, B. Acha, and R. M. Rangayyan, "Segmentation-based lossless compression for color images," in *Proc. Tenth International Conf. on Image Analysis and Processing (ICIAP'99)*, Venice, Italy, pp. 90–94 (1999).
11. G. K. Wallace, "The JPEG still picture compression standard," *Commun. ACM* **34**(4), 30–44 (1991).
12. C. S. Lim, *Two-dimensional Signal and Image Processing*, Prentice Hall Signal Processing Series, Prentice-Hall, Englewood Cliffs, NJ, (1990), pp. 422–423.
13. C. Serrano, L. Roa, and B. Acha, "Evaluation of a telemedicine platform in a burn unit," *Proc. of IEEE International Conf. on Information Technology Applications in Biomedicine*, Washington DC, pp. 121–126 (1998).
14. S. Van Assche, W. Philips, and I. Lemahieu, "Lossless compression of pre-press images using a novel color decorrelation technique," *Proc. SPIE* **3308**, 85–92 (1998).



Carmen Serrano received her MS in Telecommunication Engineering from the University of Seville, Spain, in 1996. In 1996, she joined the Signal Processing and Communication Group at the same university, where she is currently an associate professor. Her research interests concern image processing and, in particular, color image segmentation and compression, mainly with biomedical applications.



Begoña Acha Piñero studied Telecommunication Engineering at Engineering Faculty of the University of Seville from 1991 to 1996. She has been working as an associate professor since 1996 in the Signal Processing and Communications Group of the Electronic Engineering Department of the University of Seville. She is currently working toward her PhD. Her current research activities include works in the field of image processing and its medical applications.



Rangaraj M. Rangayyan received his Bachelor of Engineering degree in Electronics and Communication in 1976 from the University of Mysore at the P.E.S. College of Engineering, Mandya, Karnataka, India, and his PhD in Electrical Engineering from the Indian Institute of Science, Bangalore, Karnataka, India, in 1980. He was with the University of Manitoba, Winnipeg, Manitoba, Canada, from 1981 to 1984. He is at present a Professor with the Department of Electrical and Computer Engineering (and an Adjunct Professor of Surgery and Radiology) at the University of Calgary, Calgary, Alberta, Canada. His research interests are in the areas of digital signal and image processing, biomedical signal analysis, medical imaging and image analysis, and computer vision. His current research projects are on mammographic image enhancement and analysis for computer-aided diagnosis of breast cancer; region-based image processing; knee-joint vibration signal analysis for noninvasive diagnosis of articular cartilage pathology; and analysis of textured images by cepstral filtering and sonification. He is the winner of the 1997 Research Excellence Award of the Department of Electrical and Computer Engineering, and the 1997 Research Award of the Faculty of Engineering, University of Calgary. He was recognized by the IEEE with the award of the Third Millennium Medal in 2000, and elected as a Fellow of the IEEE in 2001.

Laura M. Roa: Biography and photograph not available.