

---

## **A novel approach to medical image compression**

---

**Matthew J. Zukoski\***

Department of Mathematics and Computer Science,  
Wilkes University, Wilkes-Barre, PA 18766, USA

E-mail: zukoski@wilkes.edu

\*Corresponding author

**Terrance Boulton**

Department of Computer Science,  
University of Colorado at Colorado Springs,  
Colorado Springs, CO 80918, USA

E-mail: tboulton@cs.uccs.edu

**Tunç Iyriboz**

College of Medicine, Penn State University,  
Hershey, PA 17033, USA

E-mail: tunc@psu.edu

**Abstract:** As medical/biological imaging facilities move towards complete film-less imaging, compression plays a key role. Although lossy compression techniques yield high compression rates, the medical community has been reluctant to adopt these methods, largely for legal reasons, and has instead relied on lossless compression techniques that yield low compression rates. The true goal is to maximise compression while maintaining clinical relevance and balancing legal risk. This paper proposes a novel model-based compression technique that makes use of clinically relevant regions as defined by radiologists. Lossless compression is used in these clinically relevant regions, and lossy compression is used everywhere else.

**Keywords:** clinically relevant; compression; lossy; lossless; JPEG; mammogram; medical imaging; mutual information maximisation; region of interest (ROI); snakes; wavelet; X-ray.

**Reference** to this paper should be made as follows: Zukoski, M.J., Boulton, T. and Iyriboz, T. (2006) 'A novel approach to medical image compression', *Int. J. Bioinformatics Research and Applications*, Vol. 2, No. 1, pp.89–103.

**Biographical notes:** Matthew J. Zukoski is completing his PhD in Computer Science at Lehigh University. He has a joint appointment as Assistant Professor in the Departments of Mathematics and Computer Science and Electrical Engineering at Wilkes University. His research interests include computer vision, medical image processing and robotics.

Terrance Boulton received his PhD in Computer Science at Columbia University in 1986, and spent eight years on the Faculty of Columbia and nine years on the Faculty of Lehigh University, where he was the founding Chairman of the Computer Science and Engineering Department. He is currently the

El Pomar Endowed Chair of Communication and Computation and Professor of Computer Science in the Department of Computer Science, University of Colorado at Colorado Springs. His research interests include computer vision, medical image processing, biometrics, surveillance and networking. He is author of more than 120 papers and he has four patents with eight patents pending.

Tunç Iyriboz is an Associate Professor of Radiology at Penn State University's Milton S. Hershey Medical Centre. His research interests include virtual colonoscopy, picture archiving and communications systems, and medical image processing.

---

## 1 Introduction

Medical image compression plays a key role as hospitals move towards filmless imaging and go completely digital. Image compression will allow Picture Archiving and Communication Systems (PACS) to reduce the file sizes on their storage requirements while maintaining relevant diagnostic information. Teleradiology sites benefit since reduced image file sizes yield reduced transmission times. Even as the capacity of storage media continues to increase, it is expected that the volume of uncompressed data produced by hospitals will exceed capacity and drive up costs.

This paper will propose an approach to improve the performance of medical image compression while satisfying both the medical team who need to use it, and the legal team who need to defend the hospital against any malpractice resulting from misdiagnosis owing to faulty compression of medical images. The improved compression performance will be accomplished by making use of clinically relevant regions as defined by physicians. Images taken of patients will be aligned to prestored image models stored in an atlas. The atlas will contain models of typical classes of images. If we are trying to compress a chest X-ray image, then it will be matched with a prestored chest X-ray model that is stored in the atlas. If we are trying to compress an X-ray of the right hand, then it will be matched with a prestored right hand X-ray model that is stored in the atlas. The atlas will be discussed in more detail later in this paper. Once an image is aligned to its corresponding model in the atlas, the two can then be aligned and the clinically relevant regions defined on this atlas image will be used to define the relevant region on the newly scanned patient image. There are varied approaches for performing the alignment such as maximisation of mutual information and deformable contour modelling.

This paper focuses on the potential gains to document the advantages of this approach. Lossless compression will be applied in the clinically relevant areas and lossy compression will be applied in the other areas. One of the reasons to use lossless compression in the relevant areas is not because radiologists think it is fine based on ROC curves. From our interviews, most doctors would prefer using lossy compression with the quality level on high. This would yield an image that is much smaller than it is lossless-compressed counterpart and the image will be 'visually lossless'. However, the lossy compressed image, when decompressed, is not identical to the original (i.e., there was loss) and the lawyers have a problem here. So this makes the problem not just what

is perceptibly 'good enough' (i.e., what is 'visually acceptable'), but rather, what is 'clinically relevant'.

In Section 2, a very brief overview of medical image compression will be given. This includes lossy and lossless techniques. In Section 3, you will find an overview of the overall approach. This section will include a discussion of how the atlas will be created, and how the alignment of images to their corresponding models will be performed. An overview of alignment will be discussed here as well. The paper ends with preliminary results using simpler regions that document the potential savings.

## **2 Medical image compression**

Before the various image compression techniques are discussed, consider the motivation behind using compression. A typical 12-bit medical X-ray may be 2048 pixels by 2560 pixels in dimension. This translates to a file size of 10,485,760 bytes. A typical 16-bit mammogram image may be 4500 pixels by 4500 pixels in dimension for a file size of 40,500,000 (40 megabytes)! This has consequences for disk storage and image transmission time. Even though disk storage has been increasing steadily, the volume of digital imagery produced by hospitals and their new filmless radiology departments has been increasing even faster. Even if there were infinite storage, there is still the problem of transmitting the images. Many hospitals have satellite centres or clinics in small towns and remote areas to make it convenient for patients who have a hard time travelling the distance to the hospital, especially for diagnostic procedures. These hospitals make use of 'teleradiology' applications that allow the clinic staff to operate the clinic without the need for a radiologist to be present. Instead of a diagnostic radiologist, a technician or basic radiologist in the clinic can take the X-ray and send the image through a network connection to the hospital where the diagnostic radiologist can read the image and send back a diagnosis. But there is a problem, especially in emergency situations where time is of the essence, because a 10 MB image will take approximately half an hour using a high-speed modem. Broadband connections such as T1 lines improve the situation, but many clinics are in such remote areas, that it is cost prohibitive to set up high-speed lines, and the 56 K modem is the most viable option. Cable modems or DSL typically have very asymmetric performance with high-speed downlinks and much slower uplinks, and the uplinks are what would limit the satellite facility. Even at standard Cable/DSL uplink speeds, a single 10 MB radiograph will take over five minutes to be sent. While that may sound reasonable, keep in mind that the patient is often asked to remain in the imaging apparatus until the radiologist has confirmed that the data is sufficient. So compression is not just about the storage costs, it is also about transmission time, imaging apparatus utilisation and convenience/comfort of the patient. Compression techniques can reduce file size and transmission time, thus improving overall care.

Image compression techniques take advantage of redundancy that occurs. There are different types of redundancy. Each compression methodology will exploit one of these redundancies. The different types of redundancies are spatial, temporal and spectral. The research presented here will focus on the first of these three types of redundancies although the techniques can be used in the others also. It will make use of spatial redundancies since static spatial X-rays will be used. These are still the most dominant type of medical imaging data used today. In future work, the issue of using temporal

redundancy will be explored as the proposed approach naturally applies on a dataset of images taken of one patient over time, and can be applied to video data as well.

Images can be compressed using lossy or lossless techniques. Lossless techniques allow the image to be compressed, then decompressed back to the original state of the image without any loss of data (Kivijarvi et al., 1998). These methods are sometimes called reversible compression methods. Compression rates for lossless techniques vary but typically are around 2:1 to 3:1. On the other hand, lossy techniques do not allow for exact recovery of the original image once it has been compressed. These methods are sometimes called irreversible compression methods. But these techniques allow for compression rates that can exceed 100:1 depending on the compress quality level and the image content. At high quality lossy levels (10:1–20:1), compression rates much greater than those obtained by lossless methods can be obtained while achieving visually indistinguishable results. That is, the human eye cannot detect a difference between the original image and the compressed-then-decompressed image with the lossy compression method. However, the medical community has been very reluctant to adopt lossy algorithms in clinical practice. This is because of the legal questions raised and the regulatory policies set by agencies such as the Food and Drug Administration. To date, there is insufficient clinical research on the use of lossy compression applied to medical images. The new compression approach, which will be proposed here utilising a hybrid lossy/lossless method, can be made all lossy or all lossless.

The most popular compression algorithms in use today in the medical community are lossless JPEG (Joint Photographic Experts Group) (Wallace, 1991) and lossless Wavelet. JPEG has been adopted by the Digital Imaging and Communications in Medicine (DICOM) group in their widely adopted DICOM image file format, but the wavelet compression algorithm is gaining ground. In fact, the DICOM Working Group added support for the JPEG 2000 standard into the DICOM format in November of 2001. It has also been adopted by ISO as a standard. JPEG 2000 is based on wavelet compression.

### *2.1 Medical compression research*

There have been numerous compression research studies examining the use of compression as applied to medical images. The papers can be categorised as focusing on just a lossless compression method, on just a lossy compression method, or focusing on both. Most have focused on lossless algorithms since the medical community has been reluctant to adopt lossy techniques owing to the legal and regulatory issues that are raised, but this situation may start to change as more lossy research is performed.

Lossless image compression is typically performed in two steps, decorrelation and coding. Image decorrelation attempts to reduce the redundancy within the image. There are several common approaches that have been taken in the literature to perform this redundancy reduction step including differential pulse code modulation, hierarchical interpolation, bit-plane encoding and multiplicative autoregression. Several popular approaches for encoding are Huffman encoding, Lempel-Ziv encoding, arithmetic encoding and run-length encoding.

Lempel-Ziv is used by Unix in the compress and gzip programs. It is also used in the GIF file format. The Huffman and Lempel-Ziv encoding methods were compared as applied to MRI images in Cohen (1991). It showed that Lempel-Ziv encoding methods achieve higher compression than compression ratios resulting from using Huffman encoding.

As mentioned earlier in this paper, lossless methods are preferred in the medical community. Of these methods, JPEG and Wavelet are most popular. These two compression methods actually gained widespread acceptance as lossy methods. However, each can be made lossless which is the preferred style in medical imaging.

The transform based lossy methods involve three stages: transformation, quantisation and lossless coding. They divide the image into subimages (for example,  $8 \times 8$  blocks in the DCT used in JPEG), and perform a transformation on the subimage, quantising and coding the resulting coefficients using some coding scheme such as Huffman, arithmetic or run-length. Discrete cosine transforms are closely related to the fourier transform and produce similar results. DCTs convert data from the spatial domain into the frequency domain.

DeVore et al. (1992) showed that the wavelet transform is a promising tool for image compression providing high rates of compression while maintaining good image quality. In 1999, we performed a quantitative comparison of three lossy compression methods (one wavelet and two JPEG) as applied to a variety of 12-bit medical images in conjunction with the Department of Radiology at the Hershey Medical Center (Hershey, PA) (Iyriboz et al., 1999). This work shows the quality of JPEG and wavelet-based compression (which is what will be used in the remainder of the study).

With regard to clinically relevant region encoding, not much has been published. In 1994, (Chen et al., 1994) made use of regions of interest using subband analysis and synthesis or volumetric datasets using wavelets. They followed up this work with (Chen et al., 1995) by using structure preserving adaptive quantisation methods as a means of improving quality for compression rates in the regions of interest. But all of their effort was on lossy approaches.

In the most relevant work, (Storm and Cosman, 1997) developed a region based coding approach. They discussed two approaches: one uses different compression methods in each region such as 'contour-texture' coding and subband decomposition coding, and the other uses the same compression method in each region such as the discrete cosine transform but with varying compression quality in each region such as by using different quantisation tables. They used two multiresolution coding schemes: wavelet zerotree coding and the S-transform, and considered only 8 bit images. In their implementation, the regions of interest were selected manually.

### **3 New approach and preliminary results**

A new model-based approach to medical image compression by the use of image registration is proposed. An image that needs to be compressed will first be aligned to an image of its own type prestored in an atlas (such as the head or chest). Once a film is registered (i.e., aligned), two possibilities exist. The simpler approach is simply to read off the 'relevant' regions and then use lossless compression in relevant regions and lossy compression in the others. The alternative is that the new image can be subtracted from the prestored atlas image generating a residual image. This residual image will be compressed (lossless in clinically relevant regions and lossy in the others). If the alignment is done well, the residual information is minimised, thus yielding higher compression.

The regions will be defined to classify areas of the image into those that are clinically relevant and those that are not clinically relevant. These regions are stored in the atlas and

have been predefined by radiologists. The areas can be considered as initial or default. It is possible that the physician currently working with the image to be compressed will feel that these default regions are inappropriate and may want to modify them. The proposed system should allow for such interaction and provide means for the physician to override the default regions and define new ones.

Lossless compression will be used in the clinically relevant regions and lossy compression will be used in areas that are not clinically relevant. Lossless and lossy compression were discussed in detail in the Introduction section. Most lossy compression algorithms such as JPEG, utilise a compression amount parameter that defines the amount of compression, and hence degradation, used on the image. Since the clinically relevant regions, by definition, delineate the regions that the physicians care about, the amount of compression performed in the clinically nonrelevant regions can be made very high. These regions will basically be used to provide a frame of reference to the physician and so long as the physician can make out in general what these clinically nonrelevant regions are, then all is well.

A prestored database of common film types such as X-rays of the head and chest can be stored. These will be used as templates and the collection will serve as the atlas. When an X-ray is taken, the software will be able to determine the gender and age of the patient based on the header information stored in the image. The header will also indicate what body part is being imaged (head, chest, foot, etc...). This information is needed to find the appropriate template in the atlas. Next, the X-ray is aligned via the proposed deformable object model matching mutual information of the image against the prestored template. Now the two images are subtracted. Whatever is left is considered a residual. The residuals are encoded by compression. If the alignment has been done well, the residuals should be minimised which means the compression has been maximised.

This image alignment model is based on a hybrid registration technique that makes use of mutual information maximisation between two images as an initial step, followed by another methodology based on deformable modelling.

Reconstruction is straightforward. The compressed residual image is first decompressed. Then the decompressed residual is added back to the template according to the model and the original image should be obtained. This would be beneficial for teleradiology applications since only the compressed residual image needs to be stored along with the atlas.

One advantage of this method is that it does not matter as much what the underlying data type is. It can be 2D or it can be 3D. Of interest is the level of compression data rate that can be achieved. It is possible with this technique to get very high compression rates with good alignment.

Another advantage of this method is that the 'residuals of residuals' can be examined for possible improvement in storage savings. Consider a person whose X-rays have been taken over time. The first X-ray can be stored after alignment. Then the difference between the second X-ray and the first X-ray after alignment is computed. This is considered the 'residual of residuals'. But there are two problems with this idea. The computing cost increases, and all of the residuals are useless if you lose the starting image. Long-term X-ray differences can also be explored.

Whether complex contours will be useful is a research study in itself, not a solution to presume. One could use an approach like MPEG, but some model that deforms smoothly would be needed, not arbitrary patches going everywhere, and but one might take a cubic spline and define a grid such that the only thing that moves on the grid when

maximisation of mutual information is performed are the nodes of the spline, and when one moves the nodes on the spline, everything else will smooth out. If one has compression in a given area of the image, the next block on the grid has to move smoothly as well from one image to the next.

The more complex issues of deformable model matching are still being studied and will be reported in later papers. The algorithm proposed here will also make use of maximisation of mutual information, (Viola, 1997).

The first level of mutual information maximisation would be the static rotation, translation and scaling, which is shown in preliminary results. More advanced techniques will first determine the scaling parameters, and then the nodes in the image will be allowed to start moving around in a subsequent pass. The global rotation, translation and scaling parameters have to be encoded as well as the deltas for the grid. Then the residuals from subtraction can be obtained.

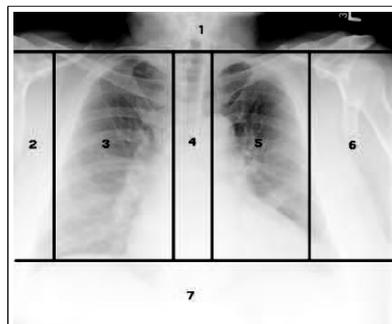
To provide supporting evidence for this ongoing study, the idea of clinically relevant regions has been applied to several medical images and the file sizes compared between the image entirely compressed lossily, and the image that used the combined lossless/lossy approach. Once the promising results were obtained, the compression system was applied to 24 X-ray images taken from male and female children and adults. Each image used in the following sections is represented at 8 bits per pixel (bpp). The results are shown below.

### 3.1 Example 1

The first image is a chest X-ray that was digitised by 'primary capture'. That is, the X-ray was captured directly from the patient using a high-resolution 12-bit digital X-ray scanner.

The image was then downsampled to 8 bits. Figure 1 shows the first example image with three clinically relevant regions defined on it. These three areas partition the image into seven areas as shown on the image indicated by regions 1–7. Regions 3 and 5 have been marked by a radiologist to be losslessly compressed. Region 4 has been marked by the radiologist to be lossy compressed, but at a higher level than the rest of the image. A JPEG level of 50 (of 100) was chosen for region 4. The clinically relevant region 3 has dimensions  $836 \times 1344$  and its upper-left coordinate is located at position 164,296. The clinically relevant region 5 has dimensions  $760 \times 1344$  and its upper-left coordinate is located at position 1288,296.

**Figure 1** Partitioned chest X-ray ( $2496 \times 2048$ , 8 bpp)



The original file size of the uncompressed raw image is 5111808 bytes using 8 bits per pixel. With the entire image compressed using lossless JPEG-2000, the compressed file size is 564185 bytes. This gives a compression ratio of 9.061: 1 or 0.883 bpp.

Regions 3 and 5 are compressed using lossless JPEG. Regions 1, 2, 6 and 7 are compressed using lossy JPEG at a compression level of 10 (out of 100). Region 4 is compressed using lossy JPEG at a compression level of 50. The results are shown in Table 1.

**Table 1** Region compressions. Overall compression rate is 0.445 bpp, almost two times smaller than pure lossless

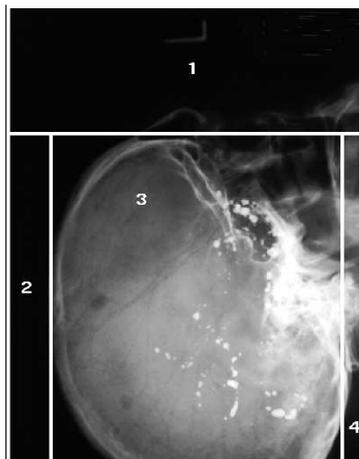
<i>Region</i>	<i>Raw bytes</i>	<i>Type</i>	<i>Compressed</i>
1	738816	Lossy	3331
2	220416	Lossy	740
3	1123584	Lossless	143308
4	387072	Lossy	3993
5	1021440	Lossless	129094
6	602112	Lossy	1478
7	1018368	Lossy	2315
Total	5111808	–	284259

So using the clinically relevant region approach, the compressed file size is 284259 bytes. This gives a compression ratio of 17.983:1 or 0.445 bpp. The clinically relevant region area takes up 42.0% of the entire image area. The compression ratio has almost doubled from the 9.061:1 ratio for the pure lossless technique.

### 3.2 Example 2

The next image is an X-ray of the skull that was digitised by ‘secondary capture’. That is, the X-ray film was digitised by using a high-resolution 12-bit scanner, and then resampled to 8 bpp. It is shown in Figure 2.

**Figure 2** Partitioned skull X-ray (1188 × 1528, 8 bpp)



The dimensions of this film are  $1188 \times 1528$ . The image is again automatically partitioned by having clinically relevant regions defined by registering to an atlas. Here is the image and the areas defined around it. Note that the clinically relevant region extends vertically to the bottom of the image. Therefore, the partitioning results in only 4 areas in this example. The clinically relevant region 1 has dimensions  $996 \times 1078$  and its upper-left coordinate is located at position 142,452.

The original file size of the uncompressed raw image is 1815264 bytes using 8 bits per pixel (bpp). With the entire image compressed using lossless JPEG-2000, the compressed file size is 769826 bytes. This gives a compression ratio of 2.358:1 or 3.393 bpp. Region 3 will be compressed using lossless JPEG. Regions 1, 2 and 4 will be compressed using lossy JPEG at level 10. The results are shown in Table 2.

**Table 2** Region Compressions. Even with a large region, overall compression rate is 32.5% better than pure lossless compression

<i>Region</i>	<i>Raw bytes</i>	<i>Type</i>	<i>Compressed (bytes)</i>
1	536976	Lossy	1647
2	152792	Lossy	490
3	1071696	Lossless	516271
4	53800	Lossy	963
Total	1815264	–	519371

So using the clinically relevant region approach, the compressed file size is 519371 bytes. This gives a compression ratio of 3.495:1 or 2.289 bpp. The clinically relevant region area takes up 59.0% of the entire image area. So, even though the clinically relevant region is so large, the compression ratio has still been improved over the 2.358:1 ratio given earlier in this example.

### 3.3 Batch study

The compression system was used to compress 24 chest X-rays. The X-rays are all digitised radiographs taken using primary capture. They range in dimension from  $1452 \times 1216$  to  $3072 \times 2560$ . The average file size is  $2500 \times 2300$ . The X-rays were taken from male and female children and adults. Each image was matched to its corresponding classic film in the atlas and aligned. After alignment, lossless compression was used in the atlas-defined regions of interest. The compression system allows additional regions of interest that are to be compressed with lossy compression but at higher quality than the remainder of the image that is compressed with lossy compression at low quality. For the chest X-rays, radiologists outlined the lung areas to be compressed losslessly, and the area between the lungs to be compressed lossily but at high quality. The remainder of the image is to be compressed lossily using low quality. The high quality area was compressed lossily at a level of 60 of 100. The areas not defined by regions of interest were compressed lossily at a level of 20 of 100.

A collage of all 24 chest X-rays is shown in Figure 3.

**Figure 3** Chest X-ray collage (8 bpp each)

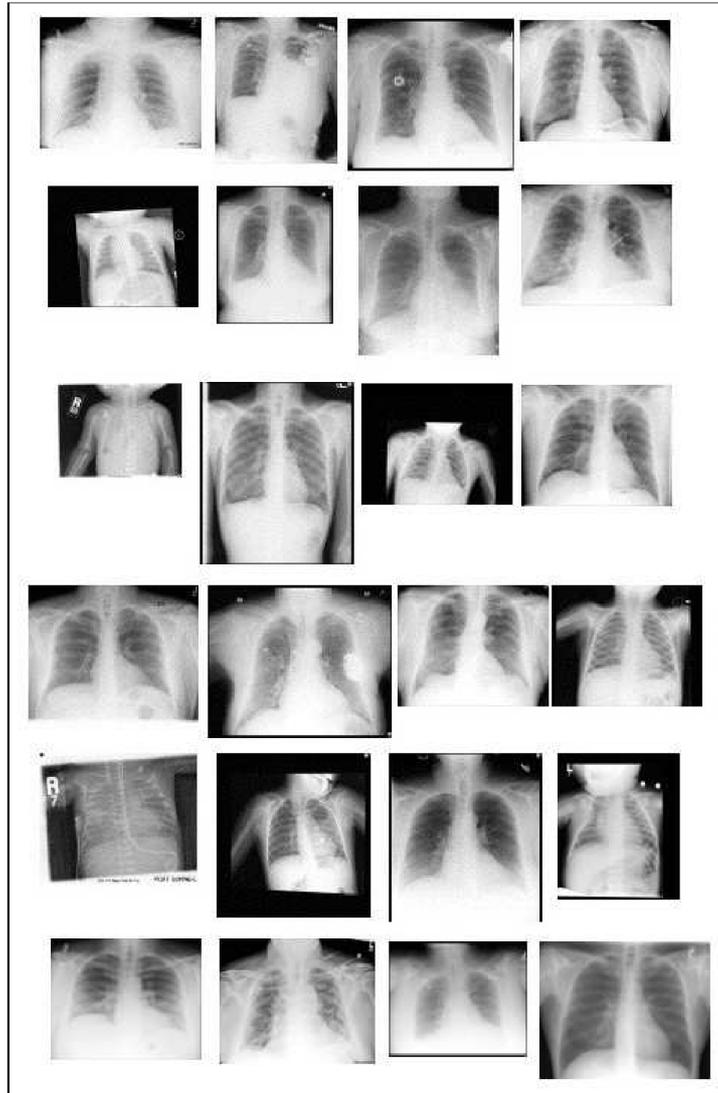


Table 3 shows the resulting compression ratio from using the lossless-only compression and hybrid compression, as well as the bits per pixel resulting from the lossless-only and from the hybrid compression technique. The last column shows the savings resulting from using the hybrid compression methodology as compared to using just lossless compression. On an average, the hybrid compressed file size was 2.65 times smaller than its corresponding lossless-only compressed image. This translates to an average file size that is 41% smaller than the lossless-only compressed file size.

**Table 3** Hybrid compression results

<i>Image no.</i>	<i>Lossless compress ratio</i>	<i>Hybrid compress ratio</i>	<i>Lossless (bpp)</i>	<i>Hybrid (bpp)</i>	<i>Hybrid savings</i>	<i>Improvement (%)</i>
1	1.7:1	3.5:1	4.7	2.3	2.1	52
2	1.7:1	5.3:1	4.6	1.5	3.0	67
3	1.7:1	3.3:1	4.7	2.4	1.9	48
4	1.7:1	2.8:1	4.8	2.8	1.7	41
5	1.7:1	8.4:1	4.6	1.0	4.9	79
6	1.6:1	3.6:1	5.1	2.3	2.3	56
7	1.5:1	2.9:1	5.4	2.7	2.0	49
8	1.7:1	3.0:1	4.7	2.6	1.8	44
9	1.5:1	4.3:1	5.3	1.8	2.9	65
10	1.6:1	4.3:1	5.0	1.9	2.7	63
11	2.1:1	11.6:1	3.9	0.7	5.6	82
12	1.6:1	3.2:1	4.9	2.5	2.0	49
13	1.6:1	4.0:1	4.9	2.0	2.5	60
14	1.6:1	3.8:1	5.0	2.1	2.4	58
15	1.8:1	3.5:1	4.5	2.3	2.0	49
16	1.8:1	6.0:1	4.5	1.3	3.3	70
17	1.5:1	4.5:1	5.2	1.8	3.0	66
18	1.7:1	5.6:1	4.8	1.4	3.3	70
19	1.8:1	4.0:1	4.5	2.0	2.2	55
20	1.9:1	6.5:1	4.2	1.2	3.4	71
21	1.8:1	4.4:1	4.6	1.8	2.5	60
22	1.6:1	3.3:1	5.0	2.4	2.1	52
23	1.5:1	3.9:1	5.2	2.1	2.5	60
24	1.5:1	2.5:1	5.3	3.2	1.7	39

### 3.4 Residual image

The previous subsections provide an overview of compression with respect to the clinically relevant regions on the raw image. An important question is if by using the atlas we can ultimately produce better compression using a ‘residual’ image. Then this residual image will be compressed in a similar region-based approach. For example, Figure 6 shows the partitioned residual image produced when the image used in Example 1 is subtracted from the image obtained (see Figure 5) after matching the corresponding atlas image (see Figure 4) to the Example 1 image. The partitioning is the same as in Example 1. As before, the three clinically relevant areas partition the image into 7 areas as shown on the image indicated by regions 1 through 7. The clinically relevant region 3 has dimensions  $836 \times 1344$  and its upper-left coordinate is located at position 164,296. The clinically relevant region 5 has dimensions  $760 \times 1344$  and its upper-left coordinate is located at position 1288,296.

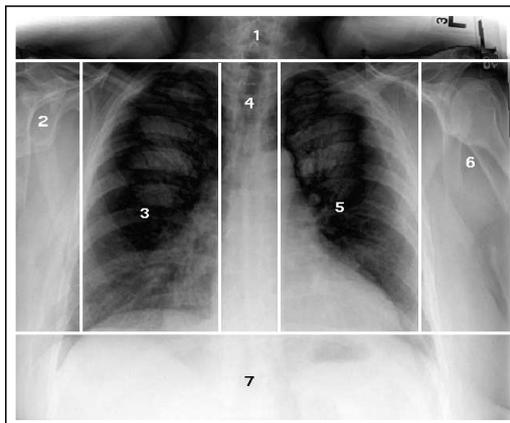
**Figure 4** Atlas chest X-ray that matches to Figure 1 (2496 × 2048, 8 bpp)



**Figure 5** The atlas image after being registered to Figure 1 (2496 × 2048, 8 bpp)



**Figure 6** Partitioned residual chest X-ray (2496 × 2048, 8 bpp)



The original file size of the uncompressed raw image is 5111808 bytes using 8 bits per pixel (bpp). With the entire image compressed using lossless JPEG-2000, the compressed file size is 2987890 bytes. This gives a compression ratio of 1.711:1 or 4.676 bpp. Regions 3 and 5 will be compressed using lossless JPEG. Regions 1, 2, 6 and 7 will be compressed using lossy JPEG at level 10. Region 4 will be compressed using lossy JPEG at level 50. The results are shown in Table 4.

**Table 4** Residual region compressions

<i>Region</i>	<i>Raw bytes</i>	<i>Type</i>	<i>Compressed</i>
1	738816	Lossy	6640
2	220416	Lossy	1362
3	1123584	Lossless	668751
4	387072	Lossy	90492
5	1021440	Lossless	606505
6	602112	Lossy	3723
7	1018368	Lossy	2756
Total	5111808	–	1380229 bytes

So using the clinically relevant region approach on the residual, the compressed file size is 1380229 bytes. This gives a compression ratio of 3.704:1 or 2.160 bpp. The clinically relevant region areas take up 42.0% of the entire image area. The compression ratio has improved over 100% over the 1.711:1 ratio for the lossless compression of the residual image.

While only a single example, it is important to note that compared to applying the clinically relevant region approach to the original image, the use of the residuals may not be more effective. In this example, residuals produced a final result of 2.160 bpp compared to the 0.445 bpp achieved using the same clinically relevant regions on the original images. There are fundamental reasons to expect that the residual may not help. The first problem is that the residual approach must encode the atlas used and ‘transform’ parameters. The more difficult problem is that even minor misalignments result in high-amplitude high-frequency data (i.e., the residual image looks mostly like edges), which are then harder to compress. This suggests that the simple 6-parameters affine transforms used for alignment are insufficient and motivates the development of more general deformable models

#### 4 Conclusion

While the residual approach is still experimental, this paper has shown that the overall approach of clinically relevant regions has clearly demonstrated advantages over both traditional lossless compression and simple lossy compression, and that the residual approach holds some promise. A larger scale study in conjunction with Penn State University’s Her-shey Medical Center is ongoing. This study will include a qualitative component to examine the effects of the hybrid compression methodology in a clinical setting. The results of this research will be published at a later date.

## References

- Chen, C.W., Zhang, Y.Q. and Parker, K.J. (1994) 'Subband analysis and synthesis of colometric medical images using wavelet', *Visual Communication and Image Processing '94*, Vol. 2306, No. 3, pp.1544–1555.
- Chen, C.W., Zhang, Y.Q., Luo, J. and Parker, K.J. (1995) 'Medical image compression with structure-preserving adaptive quantization', *Visual Communication and Image Processing '95*, Vol. 2501, No. 2, pp.983–994.
- Cohen, L.D. (1991) 'On active contour models and balloons', *Computer Vision, Graphics, and Image Processing. Image Understanding*, Vol. 53, No. 2, pp.211–218.
- DeVore, R.A., orn Jawerth, B. and Lucier, B.J. (1992) 'Image compression through wavelet transform coding', *IEEE Transactions on Information Theory*, March, pp.719–746.
- Iyriboz, T.A., Zukoski, M.J., Hopper, K.D. and Stagg, P.L. (1999) 'A comparison of wavelet and joint photographic experts group lossy compression methods applied to medical images', *Journal of Digital Imaging*, May, Vol. 12, pp.14–17.
- Kivijarvi, J., Ojala, T., Kaukoranta, T., Kuba, A., Nyul, L. and Nevalainen, O. (1998) 'A comparison of lossless compression methods for medical images', *Computerized Medical Imaging and Graphics*, No. 22, pp.323–339.
- Storm, J. and Cosman, P.C. (1997) 'Medical image compression with lossless regions of interest', *Signal Processing*, Vol. 59, No. 2, pp.155–171.
- Viola, P. (1997) 'Alignment by maximization of mutual information', *International Journal of Computer Vision*, Vol. 24, No. 2, pp.137–154.
- Wallace, G.K. (1991) 'The JPEG still picture compression standard', *Communications of the ACM*, Vol. 34, No. 4, April, pp.30–44.

## Bibliography

- Boulton, T., Fenster, S. and O'Donnell, T. (1994) 'Physics in a fantasy world vs robust statistical estimation', *Object Representation in Computer Vision*, Springer Verlag Lecture Notes in Comp. Science, Vol. 994, pp.227–296.
- Brown, L.G. (1992) 'A survey of image registration techniques', *ACM Computing Surveys*, Vol. 24, No. 4, pp.325–376.
- Van den Elsen, P.A. and Viergever, M.A. (1993) 'Medical image matching – a review with classification', *IEEE Engineering in Medicine and Biology*, Vol. 12, March, pp.26–39.
- Gokturk, S., Tomasi, C., Girod, B. and Beaulieu, C. (2001) 'Medical image compression based on region of interest, with application to colon CT images', *23rd Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, Istanbul, Turkey, Vol. 3, pp.2453–2456.
- Kass, M. and Witkin, A. and Terzopoulos, D. (1988) 'Snakes: active contour models', *IJCV*, pp.321–331.
- Kass, M., Witkin, A. and Terzopoulos, D. (1987) 'Snakes: active contour models', *Proc. of IEEE Conference on Computer Vision*, pp.259–268.
- Ke, L. and Marcellin, M.W. (1995) 'Near-lossless image compression: minimum- entropy, constrained-error DPCM', *C/P*, pp.298–301.
- National Electrical Manufacturing Association (NEMA) (2001) *The DICOM Standard*, Suite 184, 3000 North 14th St., Rosslyn, VA, USA.
- National Electrical Manufacturing Association (NEMA) (2001) *The DICOM Standard Version 3.0 Final Draft*, <http://medical.nema.org>.
- O'Donnell, T., Fang, X.S., Boulton, T.E. and Gupta, A. (1998) 'The extruded generalized cylinder: a deformable model for object recovery', in *IEEE Proceedings of the Conference on Computer Vision and Pattern Recognition*, Seattle, WA, June, pp.174–181.

- Rabbani, M. and Jones, P.W. (1991) 'Digital image compression techniques', *SPIE Optical Engineering Press*, Bellingham, Washington.
- Ratnakar, V. (1997) *Quality-Controlled Lossy Image Compression*, PhD Thesis, University of Wisconsin, Madison.
- Shen, L. and Rangayyan, R.M. (1997) 'A segmentation-based lossless image coding method for high-resolution medical image compression', *IEEE Trans. Medical Imaging*, Vol. 16, No. 3, pp.301–307.
- Tavakoli, N. (1991) 'Lossless compression of medical images', *Proceedings 4th Annual IEEE Symp. on Computer-Based Medical Systems*, pp.201–207.