An Analytical Look at the Effects of Compression on Medical Images

Kenneth Persons, Patrice Palisson, Armando Manduca, Bradley J. Erickson, and Vladimir Savcenko

This article will take an analytical look at how lossy Joint Photographic Experts Group (JPEG) and wavelet image compression techniques affect medical image content. It begins with a brief explanation of how the JPEG and wavelet algorithms work, and describes in general terms what effect they can have on image quality (removal of noise, blurring, and artifacts). It then focuses more specifically on medical image diagnostic content and explains why subtle pathologies, that may be difficult for the human eye to discern because of low contrast, are generally very well preserved by these compression algorithms. By applying a wavelet decomposition to the whole image and to specific regions of interest (ROI), and by understanding how the lossy quantization step attenuates signals in those decomposition energy subbands, much can be learned about how tolerant various anatomical structures are to compression. High-frequency anatomical structures that have their energy represented by a few large coefficients (in the wavelet domain) will be well preserved, while, those structures with high frequency energy distributed over numerous smaller coefficients are the most vulnerable to compression. Digitized films showing subtle chest nodules, a subtle stress fracture, and CT and MR images are used to show these results.

KEY WORDS: compression, wavelet compression, JPEG compression, teleradiology, PACS, medical image compression, effects of compression

THE NEED FOR reduced transmission time and the massive archive storage requirements for digital medical image data has fostered an increased interest in the use of effective compression techniques. Lossless compression algorithms allow for perfect reconstruction of the original data, but achieve compression ratios of only 2 or 3:1. Lossy techniques based on joint photographic experts group (JPEG) or wavelet compression can reach higher ratios (from 5:1 to 100:1) in exchange for a slight to moderate degradation of the image. This report will take a closer look at how compression affects medical images. We will begin with a brief explanation of how lossy image compression techniques work. We will show some of the general effects that lossy compression can have on image quality. Then we will focus more specifically on how compression affects the diagnostic content of particularly subtle pathologies, ones that you might expect would be vulnerable to compression. Finally, we will show how some types of images are more tolerant of compression than others and how spectral analysis can give us some clues that help explain why.

IMAGE COMPRESSION

Among lossy compression methods, JPEG and wavelet-based compression schemes have been widely used for medical images.[1-6] They are both based on a well defined process that involves three steps. First, an image transformation (sometimes referred to as decorrelation or signal decomposition) maps the image from greyscale values in the spatial domain to coefficients in the frequency or wavelet domain. This transformation is fully reversible, which means no information is lost during this step. The second step is a quantization stage. This is the lossy part of the algorithm, where the coefficients are approximated or truncated according to factors such as amplitude, position, and the amount of compression desired. Finally, an encoding process, generally based on an entropy coding scheme such as Huffman or arithmetic coding, losslessly compresses the remaining coefficients into a compact data stream that represents the compressed image. Decompression simply reverses the process by doing an entropic decoding followed by an inverse of the original transformation. To better understand how compression will affect medical images, we need to look a little more closely at how the JPEG and wavelet techniques approach these three steps.

JPEG Compression

The JPEG standard algorithm first decomposes the original image into 8 × 8 pixel blocks. A DCT (Discrete Cosine Transform) is applied individually to each block to generate an 8 × 8 block of coefficients representing energy in a range from

---

From the Departments of Information Services, Physiology and Biophysics, and Diagnostic Radiology, Mayo Foundation, Rochester, MN.

Address reprint requests to Kenneth Persons, 842 Centerplace, Mayo Foundation, 200 First St, SW, Rochester, MN 55905.

0897-1889/97/1003-1025$3.00/0

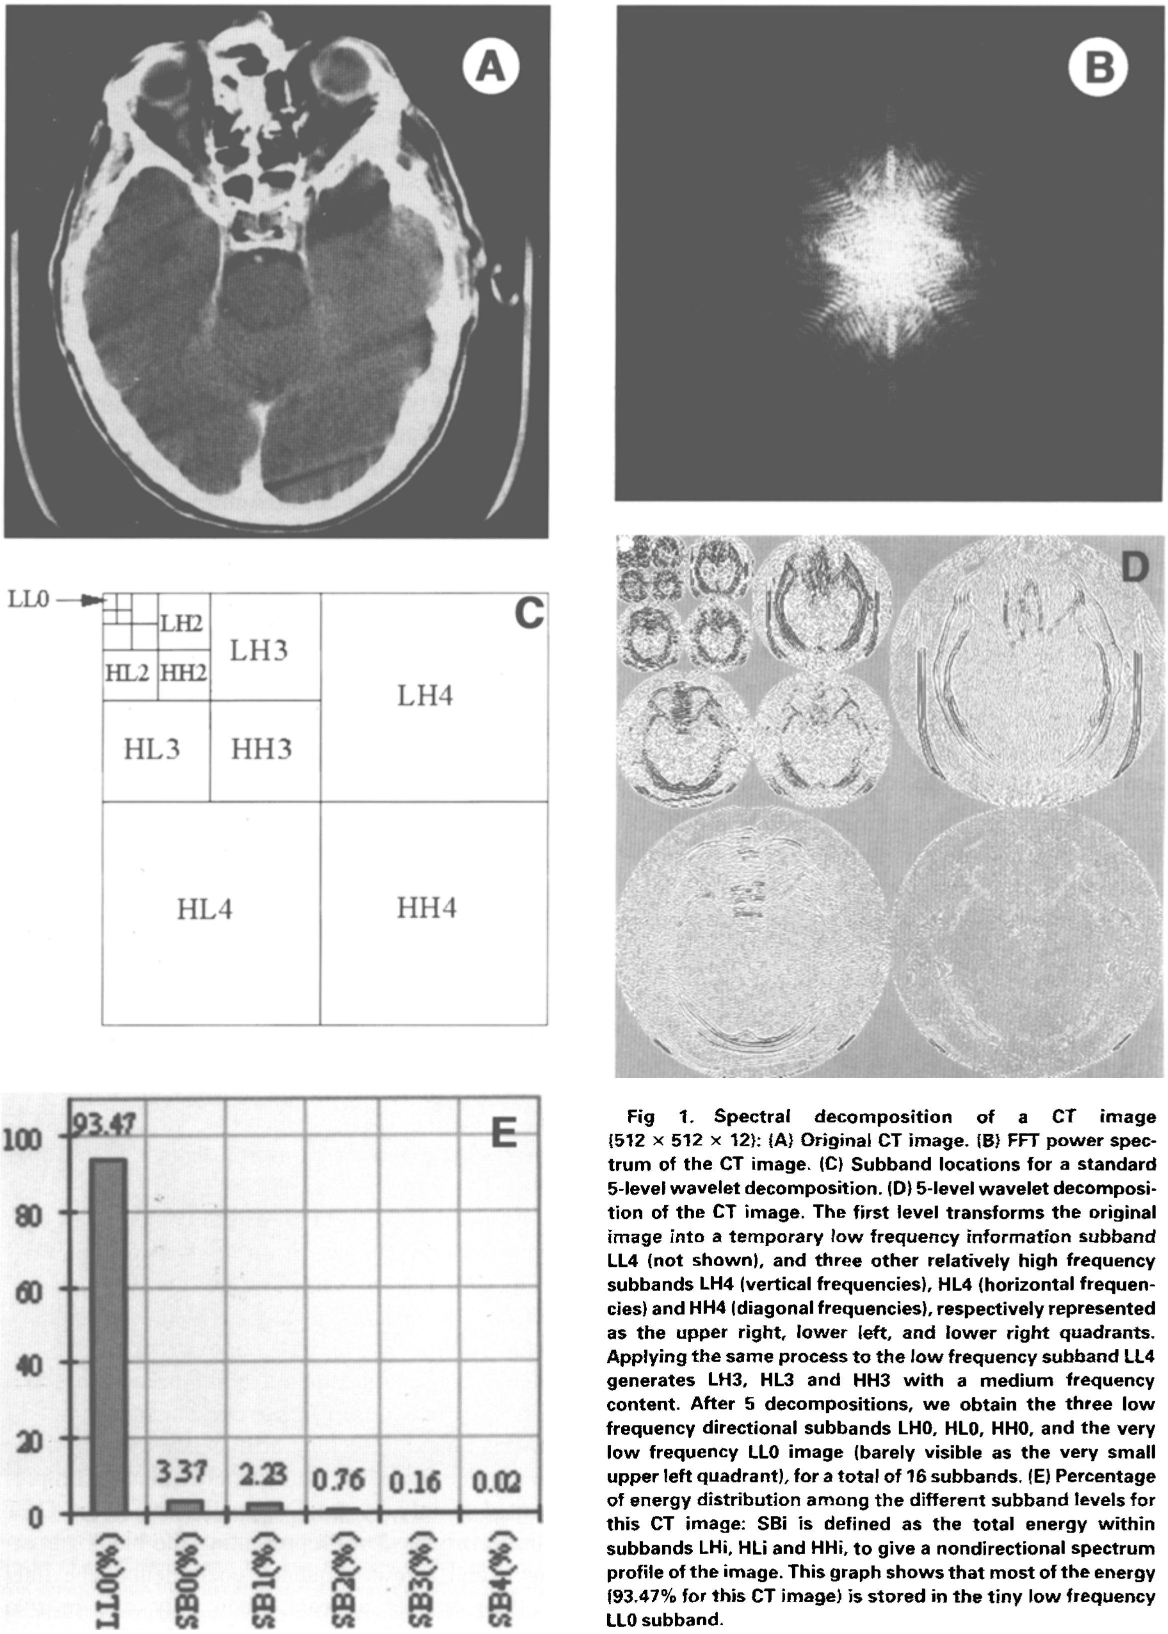

Fig 1. Spectral decomposition of a CT image (512 × 512 × 12): (A) Original CT image. (B) FFT power spectrum of the CT image. (C) Subband locations for a standard 5-level wavelet decomposition. (D) 5-level wavelet decomposition of the CT image. The first level transforms the original image into a temporary low frequency information subband LL4 (not shown), and three other relatively high frequency subbands LH4 (vertical frequencies), HL4 (horizontal frequencies) and HH4 (diagonal frequencies), respectively represented as the upper right, lower left, and lower right quadrants. Applying the same process to the low frequency subband LL4 generates LH3, HL3 and HH3 with a medium frequency content. After 5 decompositions, we obtain the three low frequency directional subbands LH0, HL0, HH0, and the very low frequency LL0 image (barely visible as the very small upper left quadrant), for a total of 16 subbands. (E) Percentage of energy distribution among the different subband levels for this CT image: SBi is defined as the total energy within subbands LHi, HLi and HHi, to give a nondirectional spectrum profile of the image. This graph shows that most of the energy (93.47% for this CT image) is stored in the tiny low frequency LL0 subband.

lower to higher frequencies. Because most of the energy in an image usually resides in the low frequency range and because the human visual system is most sensitive there, a quantization table is applied to closely preserve the low frequency coefficients, and roughly approximate the high frequency coefficients (zeroing the highest). This preserves most of the information, but significantly reduces the entropy (amount of bits needed to encode the resulting coefficients). Thereafter, a zigzag reordering of the coefficients generates long sequences of 'zeros,' which are efficiently compressed by run length encoding, while the nonzero coefficients are Huffman encoded. For our experiments, we used a standard JPEG compressor that handles the full 12 bits per pixel of greyscale information instead of scaling the original image to work with the more typical 8 bit JPEG compressors.

Wavelet Compression

The wavelet image transform, or multiscale wavelet decomposition, is usually based on a separable set of low pass/high pass filters applied several times to generate representations of the original image at various frequency scales in several orientations. The Mayo Foundation has been refining a wavelet compression scheme for the past 2 years, and unless specified, further discussion of wavelet compression in this report will refer to this particular algorithm. The wavelet compression uses a 5-level DWT (Discrete Wavelet Transform) with the 9-tap/7-tap biorthogonal filters of Antonini.[1] Such a decomposition is shown on Fig 1d for a computed tomography (CT) image of the head. The quantization and entropy encoding steps are combined using an algorithm called SPIHT (Set Partitioning in Hierarchical Trees), detailed in,[2,3] which exploits a tree-based organization that reflects the hierarchical structure of the wavelet decomposition, (ie, the relationships between coefficients representing the same location at different frequency scales), and which successfully refines coefficient values using a series of octavely decreasing thresholds. An arithmetic encoder can optionally be applied to slightly increase compression performance, with a tradeoff of a longer computation time.

The main difference in the spectral decomposition by DCT and DWT is that the DWT has its coefficients partially localized in both space and frequency, whereas the DCT coefficients are fully localized in frequency. This partial spatial localization has proven to be useful in analyzing the local spectral properties of particular regions of interest on an image such as anatomic structures or abnormalities.[6]

EFFECTS OF COMPRESSION ON IMAGE QUALITY

The effect that compression has on image quality depends on the image content, spatial and spectral distribution, and the compression level (or quality factor) which determines the degree of the quantization. The following points summarize the different effects that can be observed on images compressed with JPEG or wavelet algorithms.

Removal of noise. At low levels of compression, with both JPEG and wavelet, the quantization step mostly discards high frequency decorrelated noise whose spectral content is represented by a large number of small high frequency coefficients. This reduction in noise, without any noticeable loss of diagnostic information, makes the decompressed image more pleasing to the human eye. There is some evidence that radiologists prefer such slightly compressed/decompressed images (eg, 2K × 2.5K × 12 digitized chest films compressed at 10:1) over the original images.[7,8]

Blurring. Blurring occurs at moderate to high levels of compression with both the wavelet and JPEG algorithms when the quantization step starts to discard or very roughly approximate frequencies that describe visible structures, including coefficients that may contain useful diagnostic information.

Artifacts. At higher levels of compression, two types of artifacts can be observed with the JPEG algorithm, the "blocking" effect and "line/pattern" effect. The blocking artifact is a result of decomposition of the image into nonoverlapping 8 × 8 blocks. A separate quantization occurs for each block, which doesn't insure continuity with neighboring blocks. Blocking artifact occurs at medium to high compression ratios where blocks get represented mainly by the DC coefficient of the DCT (uniform gray level representing the block's average), and a few frequency coefficients. The line/pattern artifact appears when only one or two frequency coefficients are preserved within a block.

In this case the decompression process sees a "pure wave" of one or a few frequencies, which appears as straight lines or as a mosaic pattern bounded by the block edges when transformed back into the spatial domain. (Fig 2A). These block artifacts could be eliminated by using a full frame DCT, which computes the DCT on the whole image instead of small 8 × 8 blocks for standard JPEG, but a full frame DCT is more computationally intensive.

Blocking artifacts do not occur on wavelet compressed images because the compression is calculated on the image as a whole, but a high degree of quantization of wavelet coefficients can generate wavelet-shaped or "rice" artifacts with orientation and spatial extension depending on the subband of the most distorted coefficients. Due to the decomposition in three directions (horizontal, vertical, and diagonal), and the lesser energy usually present in high frequency bands, one is most likely to see either horizontal, vertical or diagonal "rice patterns" of short lengths in the image. As the compression ratio increases, the quantization will begin to affect lower frequency coefficients (usually with greater values), thus generating longer "rice" shaped artifacts (see example of the "rice" artifact in Fig 2B).

EFFECTS OF COMPRESSION ON DIAGNOSTIC CONTENT

Now that we have considered the general effects of compression, we will take a closer look at how compression affects the diagnostic content of medical images. Intuitively, we may expect that subtle findings (ones that are barely discernible in the original image, such as a subtle stress fracture in a bone film, or a faint nodule on a chest film) are the types of pathology that might be most vulnerable to compression. In reality, this is not always the case. Subtle pathologies, that may be difficult for the human eye to discern because of low contrast, but which have a significant spatial extent, are typically characterized by low frequencies in the spectral domain. These pathologies are quite tolerant to compression, as they are well preserved by the JPEG quantization table, or by the concentration of energy into fewer coefficients in low frequency wavelet subbands. Such subtle pathologies may remain visible even at high levels of compression. As an example, Fig 3 shows an enhanced region of interest (ROI) taken from two digitized chest films. At the center of each ROI is a small, uncalcified lung nodule, one benign and one malignant, shown as original and compressed at 40:1 and 80:1 with

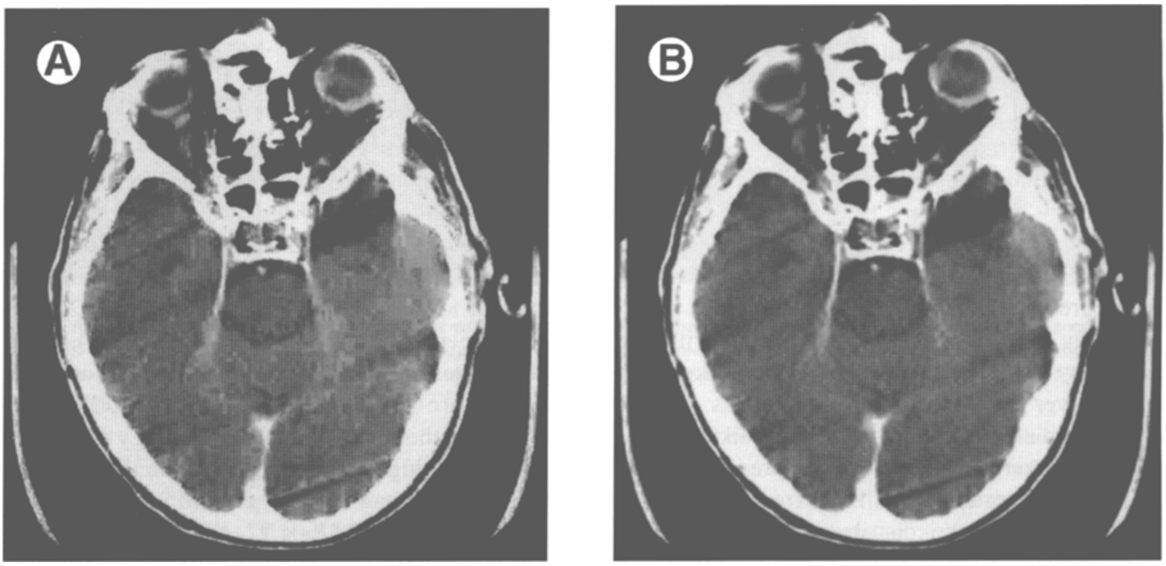

Fig 2. JPEG and wavelet compression artifact. Effects of quantization with JPEG and wavelet compression at 30:1 of the CT image in Fig 1. (A) The two main JPEG artifacts are clearly visible here: the "blocky" effect due to over quantization of the 8 × 8 blocks of coefficients, and the line artifacts within blocks. (B) Wavelet artifacts that look like "grains of rice" appear due to over-quantization that discards some wavelet coefficients and not others. Note: Our wavelet compression scheme has an option to optimize the compression quality of CT images, which almost completely eliminates the rice effect at 30:1.[4] That option was turned off for this example to emphasize the quantization effects.

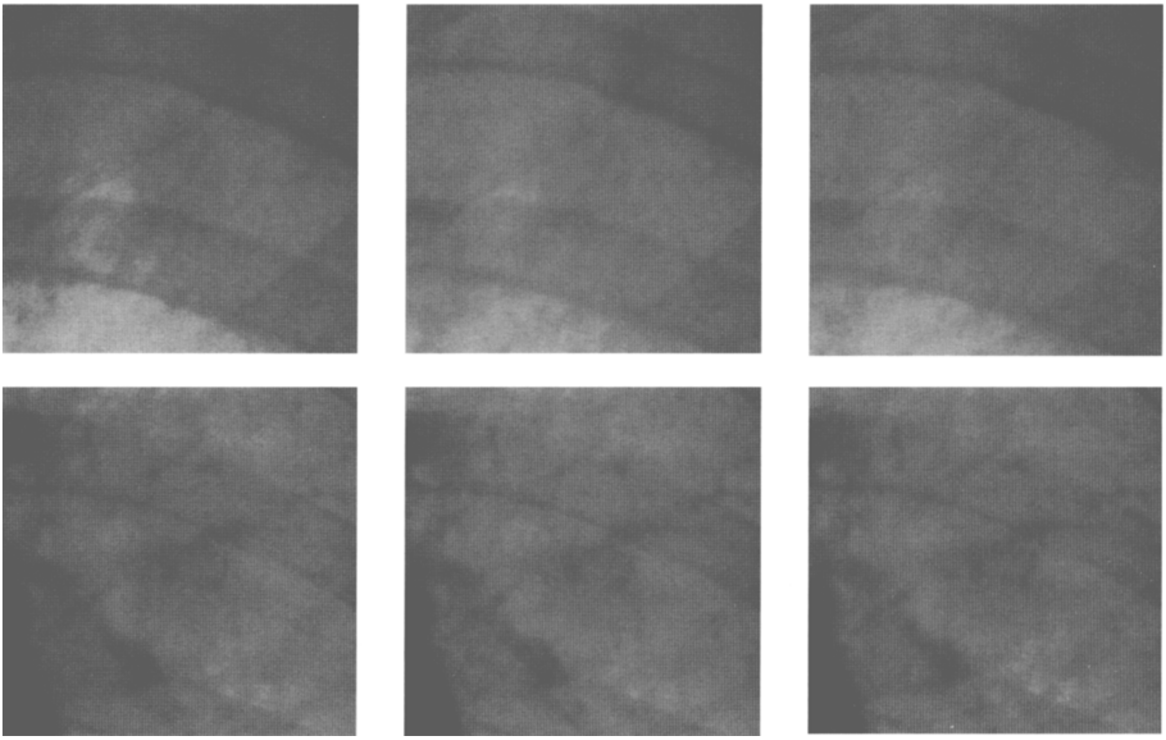

Fig 3. Subtle nodules present on two digitized chest films scanned at a 2k × 2.5k × 12 bit resolution, after magnification and contrast enhancement, and shown at wavelet compression levels of 1:1, 40:1, and 80:1. From left to right for each nodule: original, compressed at 40:1, and compressed at 80:1. The upper row shows a benign nodule, and the lower row shows a malignant nodule. The shape and contour are very well preserved even at 80:1.

our wavelet algorithm. The low contrast detail has been well preserved and the nodule shape and contour are clearly identifiable even at 80:1 compression.

It is high frequency features that are usually more vulnerable to compression. An important determining factor is how the energy is distributed among high frequency coefficients in the spectral or wavelet domain. The quantization process will better preserve high frequency pathologies represented by a few large coefficients than it will high frequency pathologies with the same energy, but spread over numerous small coefficients. This is because small coefficients are more likely to be rounded to zero, even at low compression levels. The extreme example of high frequency image content with energy distributed over numerous smaller coefficients is random noise, and this is usually discarded first, as noted above. Fine, irregular texture patterns would also contain many small, high frequency coefficients, so we would expect them to degrade easily. Such an example is shown in Fig 4 where the trabecular pattern of bone (high frequency) degrades long before a subtle fracture (lower frequency).

COMPRESSION TOLERANCE

Different types of images exhibit different degrees of *tolerance* to compression, where *tolerance* may be defined as *the range of compression where the decompressed image is acceptable for interpretation.* Subjectively, it is clear that chest films are tolerant to fairly high compression ratios (20:1 to 40:1, or even 80:1 as in the example above), while CT images are much harder to compress, and MR images are harder yet. This observation can be related to the relative amount of energy present in low versus high frequency subbands. For a set of ten typical images from each of these sources (digitized 2K × 2.5K × 12 chest films, 512 × 512 × 12 direct captured CT and 256 × 256 × 12 direct captured MR) we found that chest films averaged 99.69% of their energy in the lowest frequency (LL0) subband, versus 92.12% for CT and 78.03% for MR (see Fig 5A). Conversely, these chest films had only .31% of their energy in all of

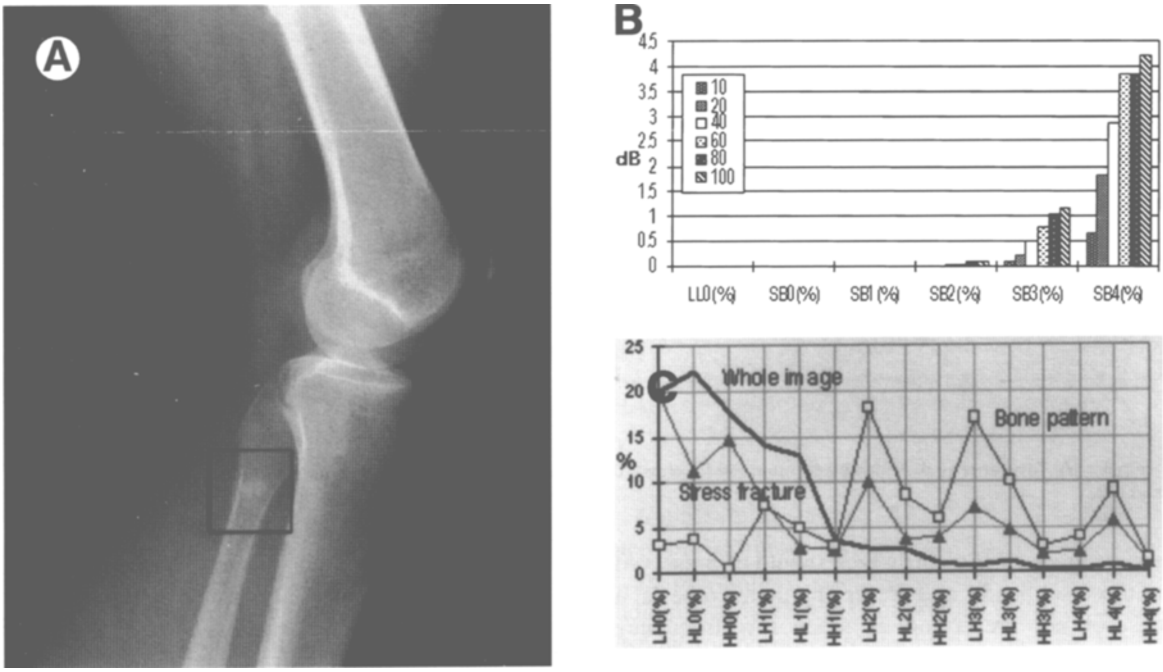

Fig 4. Energy distribution and attenuation for two ROIs in a digitized film of a subtle stress fracture. (A) Digitized film of a stress fracture (2343 rows × 1856 columns × 12 bits). The stress fracture is characterized by a focal sclerosis, inducing a slight change of density in the bone, typically low frequency. The trabecular bone pattern visible within the bone shows relatively high directional frequencies along the vertical axis of the bone. (B) Shows the attenuation of subband relative energy on the whole stress fracture image for wavelet compression ratios of 10:1 to 100:1. Notice how the quantization process focuses on removal of energy in the higher frequency subbands. (C) Shows the subband energy distribution for the whole image and two ROIs: the bone stress fracture (calcification) with low frequency contents, and the trabecular bone pattern characterized by high vertical frequencies. The higher frequency bone pattern is more vulnerable than the stress fracture to being attenuated.

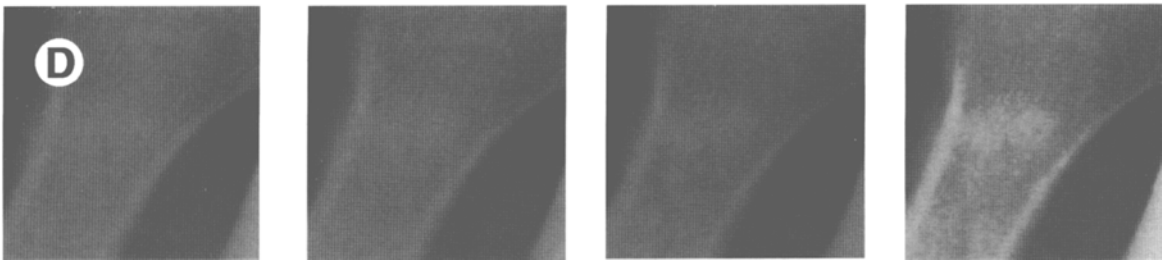

Fig 4 (cont'd). (D) show the stress fracture area with contrast enhancement: original (far left), compressed at 20:1 (second from left), 40:1 (third from left), 100:1 (fourth from left). Upper square delimits a typical bone trabecular pattern that starts getting blurred at compression ratios between 30:1 and 50:1. Lower square represents the stress fracture, which remains well preserved even when compressed at 100:1.

the other bands, compared with 7.88% for CT and 21.97% for MR (see Fig 5B). This significant high frequency energy in CT and MR images is what makes them hard to compress. We suggest that this single measure–percentage of energy in (or not in) the lowest frequency subband–is a good predictor of overall tolerance to compression for images in general (although how specific features within an image respond to compression requires a more careful analysis, as shown above). In the extreme case, an image with no high-frequency information whatsoever is oversampled, and can be compressed with no loss of information by decimation. Typical high resolution chest films appear to be close to this limit.

A related factor that affects image tolerance to compression is how the non LL0 energy is distributed in the other subbands. Sharp peaks indicate

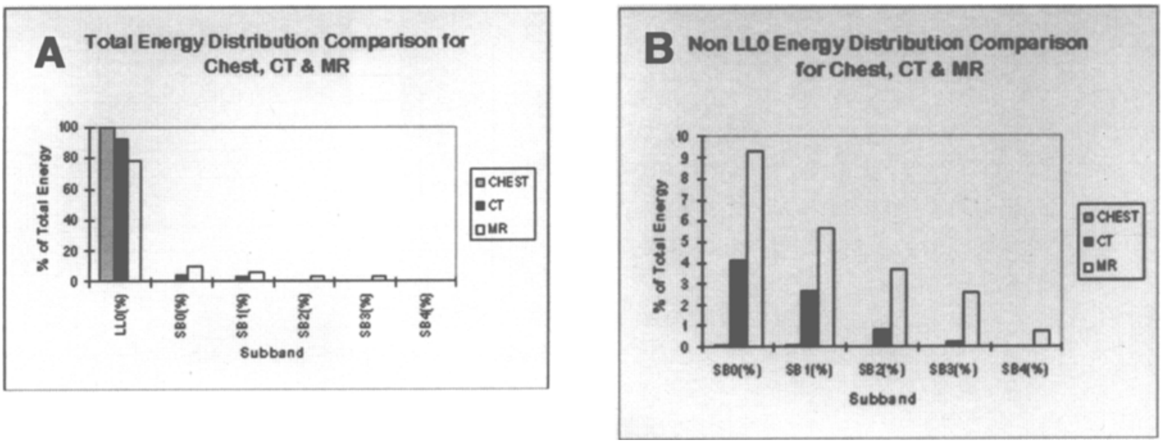

Fig 5. (A) Total energy distribution comparison for chest, CT, and MR. (B) Non LLO energy distribution comparison for chest, CT, and MR.

some higher coefficients which should be preserved in those subbands, and preserving them comes at a cost of not preserving as much low frequency information.[4] In the case of the CT image in Fig. 1, LH0 and LH1 (high frequency vertical information) contained higher peaks (not shown here) because of the contribution by the patient's head brace. Other sharp contrast information, such as text burned into the image, contributes high frequency energy that reduces the images' tolerance for compression.

CONCLUSION

Lossy compression in medical imaging naturally raises the question of whether or not clinically important information has been compromised. How much compression is acceptable? A major challenge is to find a reliable way to quantify this degradation in terms that allow us to answer this question. Objective quantitative measurements, although easy to obtain, only show partial correlation with subjective visual evaluation for diagnostic purposes. In this study, we have tried to show and explain how JPEG and wavelet algorithms can affect image quality. We have also shown that subtle low contrast pathologies that are sometimes difficult to perceive with the human eye can be quite well preserved by these compression methods. Conversely, irregular high frequency patterns are easily degraded. Thus, pathologies vary in vulnerability to compression based on how their energies are distributed in the spectral domain. The large variability of tolerance of medical images to compression, and what we have learned in this study, make it clear that further studies to evaluate compression for diagnostic purposes should focus on particular modalities and specific findings, with quantitative measurements localized in both space and frequency.

REFERENCES

1. Antonini M, Barlaud M, Mathieu P, et al: Image coding using wavelet transform. IEEE Trans Image Proc 1:205-220, 1992

2. Manduca A, Said A: Wavelet Compression of Medical Images with Set Partitioning in Hierarchial Trees. Medical Imaging Image Display, Proc SPIE 2707:192-200, 1996

3. Said A, Pearlman WA: A new fast and efficient image codec based on set partitioning in hierarchial trees. IEEE Trans Circuits and Systems for Video Tech 6:243-250, 1996

4. Manduca A, Erickson BJ, Persons K, et al: Histogram transformation for improved compression of CT images. Medical Imaging 1997: Image display, SPIE 3031 (in press)

5. Good WF, Maitz GS, Gur D: Joint Photographic Experts Group (JPEG) Compatible Data of Mammograms

6. Karson TH, Chandra S, Morehead AJ: JPEG Compression of Digital Echocardiographic Images: Impact on Image Quality. J Am Soc Echocardiogr 8:306-318, 1995

7. Cox GG, Cook LT, Insana MF: The effects of lossy compression on the detection of subtle pulmonary nodules. Med Phys 23:127-132, 1996

8. Erickson BJ, Manduca A, Persons K, et al: Evaluation of irreversible compression of digitized PA chest radiographs. J Digit Imaging 1997 (in press)