

Irreversible Compression of Medical Images

Bradley J. Erickson M.D.

The volume of data from medical imaging is growing at exponential rates, matching or exceeding the decline in the costs of digital data storage. While methods to reversibly compress image data do exist, current methods only achieve modest reductions in storage requirements. Irreversible compression can achieve substantially higher compression ratios without perceptible image degradation. These techniques are routinely applied in teleradiology, and often in Picture Archiving and Communications Systems. The practicing radiologist needs to understand how these compression techniques work and the nature of the degradation that occurs in order to optimize their medical practice. This paper describes the technology and artifacts commonly used in irreversible compression of medical images.

KEY WORDS: data compression, wavelets, JPEG

THE STORAGE and image transfer requirements of medical images have hampered attempts to implement picture archiving and communications systems (PACS) and teleradiology. Image compression recently has been explored as a means of reducing costs of managing large image data sets. Lossless compression methods use redundancy within an image to more efficiently transmit image information while allowing perfect reconstruction, but these methods achieve only 2:1 to 4:1 reduction for medical image.¹ Irreversible or "lossy" techniques can reduce images by arbitrarily large ratios, but do not perfectly reproduce the original image. However, the reproduction may be good enough that there is no perceptible image degradation nor compromised diagnostic value. This report reviews the application of image compression techniques to medical imagery, focusing on the irreversible methods, including the JPEG2000 standard. Following that is a review of measures for evaluating compression algorithm performance and some of the recent results for wavelet compression.

COMPRESSION TECHNIQUES

Most irreversible image compression techniques involve 3 steps: transformation, quantization, and encoding. Transformation is a lossless step in which the image is transformed from grayscale values in the spatial domain to coefficients in some other domain. One familiar transform is the Fourier transform used in reconstructing magnetic resonance images (MRI). Other transforms such as the discrete cosine transform (DCT) and discrete wavelet transform (DWT) are more commonly used for image compression. No loss of information occurs in the transformation step. Quantization is the step in which data integrity is lost. It attempts to minimize information loss by preferentially preserving the most important coefficients, whereas less important coefficients are roughly approximated, often as zero. Quantization may be as simple as converting floating point values to integer values. Finally, these quantized coefficients are encoded. This also is a lossless step in which the quantized coefficients are compactly represented for efficient storage or transmission of the image.

JPEG COMPRESSION

The JPEG (Joint Photographic Experts Group) compression standard is a widely used

From the Department of Radiology, Mayo Foundation, Rochester, MN.

Correspondence to: B. J. Erickson, M.D., Department of Radiology, Mayo Foundation, 200 First St. SW, Rochester, MN 55905; tel: 507-284-8548; fax: 507-284-2405; e-mail: bje@mayo.edu

Copyright © 2002 by SCAR (Society for Computer Applications in Radiology)

Online publication 30 April 2002

doi:10.1007/s10278-002-0001-z

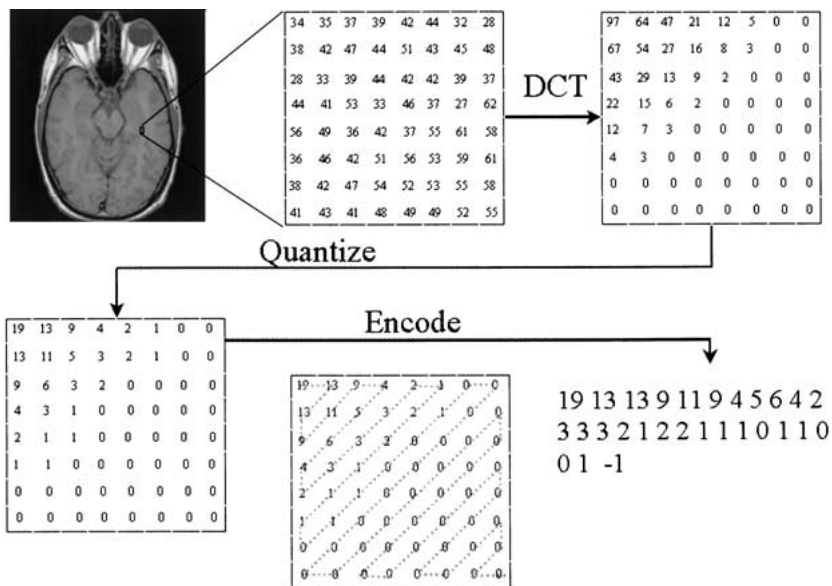


Fig 1. The JPEG Algorithm. The image is first separated into 8×8 pixel subimages. The DCT of each subimage then is computed. These coefficients then are quantized using a quantization table (for this illustration, each value is divided by 5). Finally, the quantized values are encoded from the upper left corner, with a 'marker value' sent when there are no more nonzero values.

compression method that includes both reversible and irreversible techniques, and has been described in detail by Wallace.² Although JPEG was not designed for medical imagery (ie, it was not defined for 12- or 16-bit intensity scales), it has been adapted for radiologic images as described by Gillespy and Rowberg.³ Figure 1 shows how the algorithm operates. It begins by dividing the image into 8 pixel \times 8 pixel blocks. The DCT of each image block is computed, resulting in an 8×8 block of spectral coefficients. Most of the information is concentrated in relatively few coefficients in the upper left corner of this DCT image.

Quantization is performed next. In this step, the coefficients are approximated to values that are easy to represent in a small amount of space. There is an 8×8 table (called the quantization table), which contains the values by which corresponding coefficients are to be divided. By using different values, spectral frequencies that are more important to the visual system can be preserved preferentially over less-important frequencies. The resulting values then are rounded off to the nearest integer.

JPEG encodes the quantized coefficients by reordering them in a zigzag pattern. This places the largest values first, with long strings of zeros at the end, which can be efficiently represented.

BASIC WAVELET COMPRESSION

Although the JPEG lossy algorithm is good for many types of images, it has some drawbacks when applied to radiographic images. It degrades ungracefully at high compression ratios, with prominent artifacts at block boundaries, and it cannot take advantage of patterns larger than the 8×8 pixel blocks. Wavelet-based compression schemes generally outperform JPEG in terms of image quality at a given compression ratio, and the improvement can be dramatic at high compression ratios.⁴

The DWT of an image is computed⁵ (Fig 2) using a pair of high- and low-pass filters with special mathematical properties. Many such "wavelet" filters exist, but many groups have adopted the 9-tap/7-tap bi-orthogonal filters of Antonini et al,⁵ because they seem to work well in real-world application.⁶ The 2 filters split the image into 2 components or subbands in each direction (each is half the original size). This produces 4 subband images, 1 containing the low-frequency information, 1 each for the high-frequency information in the X or Y direction, and 1 for high-frequency information in both X and Y. The process is repeated on the low-frequency component, breaking it up into "high-low" and "low-low" components. If this




Fig 2. 2a. T₁-weighted axial image of the head. 2b. Five-level DWT of this image. There is no difference in information between the two—only how it is represented.

process is performed n times, an n -level discrete wavelet transform is created. A 5-level discrete wavelet transform of an MRI is shown in Fig 2. The DWT is effective for compression because it effectively concentrates the information into a few coefficients, with most other coefficients being zero or close enough to zero that they can be considered zero without degrading the image.

Most wavelet compression algorithms compute a 4- or 5-level DWT, quantize the resulting coefficients, and efficiently encode the quantized coefficients. The quantization is performed by dividing each coefficient by a quantization parameter and rounding off to the nearest integer. Having a larger quantization parameter will result in more coefficients that are zero, and hence, increases the compression ratio. Finally, encoding converts the coefficients into values that can be stored or transmitted efficiently.

ADVANCED WAVELET TECHNIQUES

It is the way that the nonzero coefficients are encoded that differentiates the advanced wavelet compression techniques. Figure 2 graphically shows the hierarchical structure of the DWT; advanced techniques capitalize on this tree-based organization of the coefficients. The most well known of these techniques is embedded zerotree coding, described by Shapiro,⁷ and enhanced by Said and Pearlman.^{8,9} The latter approach, termed set partitioning in hierarchical trees (SPIHT), was one of the early

successful advanced wavelet techniques—it yielded significantly better results than conventional wavelet compression with similar computational complexity.⁹ In addition to resulting in efficient compression, it also transmitted the compressed bitstream in which approximations of the most important coefficients (regardless of location) are transmitted first. The values of these coefficients are progressively refined, and the most important remaining information—that which yields the largest distortion reductions—is transmitted next. It can be shown that such a transmission scheme (with uniform weighting) is the optimal way to decrease the root-mean-square (RMS) error in the reconstructed image.⁹

JPEG2000

Because JPEG was specified for computers that existed over a decade ago, and because new technologies like wavelet had surpassed JPEG for many types of images, the JPEG group set out to update the standard, which is now known as JPEG2000. For this paper, *JPEG* will refer to the older compression method, *wavelet* will refer to the family of specific wavelet methods, and *JPEG2000* will refer to the developing standard.

The JPEG2000 effort has been substantial. This group identified a number of shortcomings of the JPEG standard that JPEG2000 would address. Among these were:

1. Better performance at high compression ratios
2. A single codestream that would support irreversible and lossless compression
3. Support for many types of images (specifically including 16-bit medical images)
4. Support for many different environments (eg, high performance local area network or low-speed wide area network).

The algorithms included in JPEG2000 include the best wavelet methods and provides for flexibility in the filters used and the wavelet transform. It is radically different in the way it encodes information to allow seamless transition from irreversible to lossless image transmission. It allows applications to apply different compression ratios to different portions of an image. Finally, it provides mechanisms for user-specifiable pixel accuracy. Not all of these features exist in the first rollout of JPEG2000. The first step only supports wavelet image encoding. The more advanced features will be finalized as later steps. The interested reader is directed to the JPEG2000 web page for details—the address is: <http://www.jpeg.org/JPEG2000.htm>.

EFFECTS OF IRREVERSIBLE COMPRESSION ON IMAGES

The alteration in an image that has been irreversibly compressed depends heavily on the characteristics of the image, the compression algorithm, and the compression ratio being used. When low compression rates are used, the quantization step largely discards high-frequency noise in which spectral content is represented by a large number of low magnitude coefficients.¹⁰ At these very low compression ratios, image degradation is imperceptible (referred to as “visually lossless”; Fig 3a). As noted above, JPEG2000 allows users to specify the maximum change in pixel value permitted. Setting this value in the context of the viewing conditions can guarantee that the alteration always will be within an acceptable range.

As the compression ratio is increased, the first perceptible changes typically are removal of “salt-and-pepper” noise (obviously, this depends on the image, but generally is true for


medical images). Because it is still largely noise that has been discarded,¹⁰ no structure can be seen in the error image; these denoised images generally are preferred by observers¹¹ and actually may improve diagnostic accuracy.¹² These differences are most easily observed on a computer display by rapidly switching between the original and compressed image,¹³ or by subtracting the compressed image from the original image.

At moderate levels of compression, blurring can be seen as the quantization step more roughly approximates coefficients that describe important features. At this point, recognizable features will be seen in subtraction images.

At still higher levels of compression, blurring increases, and artifacts that are characteristic for a particular algorithm appear. Two types of artifacts can be observed with the JPEG algorithm, the “blocking” effect and “line-pattern” effect. Both result from decomposing the image into nonoverlapping 8×8 blocks and quantizing each block separately.¹⁴ Blocking artifacts do not occur on wavelet compressed images, because the transformation and quantization are calculated on the image as a whole rather than on small blocks, but a high degree of quantization of wavelet coefficients can generate wavelet or “rice-shaped” artifacts with orientation and spatial extension that correspond with the subband of the most distorted coefficients.¹⁴ Figure 4 shows examples of the types of artifacts that are characteristic of the JPEG and wavelet algorithms at extreme compression levels.

The great concern in using irreversible compression for medical images is that subtle findings (eg, a faint nodule on a chest film) would be “lost” in the compressed image, but this is not always the case. Subtle findings may be difficult for the human eye to discern because of low contrast, but if they have a significant spatial extent, they are characterized by low frequencies in the spectral domain, which are well preserved by wavelet compression (and most other compression schemes). Such subtle findings may remain visible even at high levels of compression.¹⁴

Features that have their energy spread over numerous smaller coefficients in the wavelet or spectral domain are most vulnerable to compression for most irreversible compression



(a) 10:1

(b) 50:1

(c) 200:1

Fig 3. In this figure, the top row of images shows chest radiographs compressed (using SPIHT wavelet method) at 10:1, 50:1, and 200:1. A magnified subregion is shown in the upper right corner. The absolute error image, using a display

width equal to 1% of the dynamic range of the image was used. Notice that at 10:1, the error image is noise. At 50:1 and 200:1, noticeable features can be seen in the error image.

methods. An example of image content with energy distributed over numerous smaller coefficients is random noise, and, as noted earlier, this usually is discarded first. Fine, irregular textures also contain many small, high-frequency coefficients and tend to exhibit blurring at moderate levels of compression. Examples include (1) white matter in a brain computed tomography (CT) image, (2) the trabecular pattern of bone radiographs, and (3) speckle in an ultrasound (US) image. These structural textures are good indicators of when compression is introducing visible loss to the data.

It also is important to recognize that some types of images tolerate much higher levels of compression than others, in which “compression

tolerance” is defined as the maximum compression in which the decompressed image is acceptable for interpretation and aesthetics. Digitized chest radiographs are very tolerant of compression (at least 40:1 for SPIHT wavelet¹¹), digitized bone films are moderately tolerant (between 20:1 and 40:1), and CT, MRI, and US images exhibit fairly low tolerance to compression (less than 20:1). Unfortunately, one cannot assign a single compression ratio for a modality even for a given organ system. In a recent study¹⁵ applying JPEG compression to a large number of head CTs and head MRIs, there were images considered acceptable by 5 of 5 viewers at ratios as high as 22, whereas other images were considered unacceptable at a ratio of 5.3.




Fig 4. Examples of the types of artifacts that are characteristic of the JPEG and wavelet algorithms at high compression levels. The left image, 4a, is the original CT subregion. 4b. Same subregion compressed at 40:1 using wavelet. Notice the “rice-shape” artifacts in the image. These

are wavelet-shaped artifacts that result from inaccuracies in coefficients at this extreme compression ratio. 4c. Same subregion also at 40:1 with JPEG. Note the clear blocks in the image. These are the 8×8 pixel subimages that are compressed individually in the JPEG algorithm.

EVALUATION OF IRREVERSIBLE COMPRESSION

Because a single compression ratio cannot be broadly applied to all modalities, or even to all images of a single modality, the need for careful evaluation of appropriate compression ratios becomes apparent. Three categories of methods have been used. (1) numerical analysis of pixel values before and after compression, (2) subjective observer evaluation focusing on aesthetic acceptability and estimated diagnostic value, and (3) objective measurement of diagnostic accuracy using blinded evaluation methods. For both subjective and objective evaluations, it is important to evaluate both low- and high-frequency structures or pathology. We will refer to this type of approach as a *dual-frequency* or *dual-pathology* method.

NUMERICAL EFFECTS OF IRREVERSIBLE COMPRESSION

The most basic measure of compression fidelity is to compute the mean pixel error for the compressed image. This is both familiar and simple but fails to measure local degradations that can lead to loss of important information.¹⁶ Attempts to correlate numerical measures with observer ratings or performance has borne little fruit.^{15,17–20} In the study by Erickson,¹⁵ 200 CT and 200 MR images of the head compressed at two different quality factors were rated by 5

neuroradiologists (4,000 total ratings), none of 47 measures correlated to any significant degree with observer ratings. Although there is a general correlation, there is no consistent threshold value that corresponds with the point at which image degradation either becomes perceptible or aesthetically unacceptable. Similarly, there is little correlation between these measures and the degradation of diagnostic quality as measured by diagnostic tasks.

SUBJECTIVE ASSESSMENT OF COMPRESSION

Many different methods of using subjective observer perceptions of images exist to evaluate the effect of compression. Many of the early studies (for example Ishiyaki et al²¹) used rankings: if the observer correctly ordered the images from least compressed to most compressed, then definite differences are presumed to exist at each step. However, it was difficult to interpret cases in which rankings were out of order. Others used subjective ratings of the appearance of a pathologic process (eg, liver masses²² or multiple sclerosis lesions²³). A weakness of this methodology is that one pathologic process may not have features susceptible to the compression method used, and the appearance of pathologic processes can be unpredictable, making it difficult for observers to determine what is degradation and what is variation in appearance. Others have used “image processing experts” to define a “just




Fig 5. Comparison of increases in disk density (thick solid line), speed of wide area networks (light solid line), the number of bytes of image data produced by CT (dashed line), and MR scanners (dotted line) in a working day. This shows

how improvements in disk density and network speeds have been matched by the radiology department's ability to create image data, and suggests that waiting for technology improvements to reduce costs will not be a successful strategy.

noticeable difference" for selecting the point at which compression resulted in a detectable difference in a group of mammograms.²⁴

One of the more popular and robust methods is to use a double-blinded 2-alternative forced-choice assessment.²⁵⁻²⁸ Using this technique, one study had chest radiologists rate the appearance of several normal anatomic structures seen on chest radiographs¹¹ and found no difference in the subjective quality of any structure up to a compression ratio of 40:1, and only a slight preference for uncompressed images on 1 of the 11 structures (vertebral body interspaces) at 80:1. Furthermore, a slight preference for compression at low ratios over originals was noted. It is suspected that this is because of the filtering properties of wavelet transformation eliminating noise.¹⁰

Others have suggested an unblended 2-alternative choice, allowing the observers to rapidly switch between images on a computer screen. The observer then is asked to determine whether there is a difference between the original and compressed image. This method recently was applied to a variety of medical images using several different JPEG2000 methods and the original JPEG algorithm.²⁹ They found that the proposed JPEG 2000 scheme appears to offer similar or improved image quality performance relative to the current JPEG standard for compression of medical images, except for radiographs, in which the performance for JPEG2000 was better. Savcenko et al³⁰ also compared JPEG2000 with the

original JPEG lossy compression in a large series of CT and MR images of the head. That study also found that JPEG and JPEG2000 were not substantially different in the maximum "acceptable" compression ratio.

DIAGNOSTIC EVALUATIONS

Most people agree that one can perceive "changes" in the image long before an image is degraded enough to lose its diagnostic value,^{1,31,32} but this cannot yet be assumed for all compression methods, modalities, or diagnostic tasks. Because the role of the radiologist usually is to make a diagnosis, carefully designed studies that measure the effects on clinical practice are essential. Many such studies have evaluated the diagnostic accuracy of compressed images.^{12,28,31,33-35} Unfortunately, they are difficult to compare because different algorithms, image types, and evaluation paradigms were used. Although some types of compression may excel for certain features (eg, maintaining low frequency contrast) they may do less well with other features (eg, high-frequency edges). Therefore, it is the author's opinion that receiver-operator characteristic (ROC) studies evaluating both low- and high-frequency features as well as textures are likely to be most valuable. Further, many studies use the original images as the gold standard, which is biased against compression because any difference can only favor uncompressed images.

An independent gold standard should be used whenever possible.²⁵

In general, JPEG does not achieve compression ratios as high as most wavelet methods. The maximum ratio for JPEG is typically between 10:1 and 20:1 for most modalities,^{23,24,36-38} perhaps because breaking the image into small blocks decreases its ability to take advantage of low-frequency features larger than 8 pixels. The logical improvement over JPEG is to stop breaking the image into small blocks and perform the discrete cosine transform over the entire image. For this "full-frame" discrete cosine transform method, compression ratios of about 20:1 have been reported to have no diagnostic degradation for chest radiographs^{33,34,39} with lower ratios for mammography.⁴⁰ It is important to note, however, that although no statistically significant difference was detected, the trend was for JPEG to result in some image degradation at these ratios, and it is possible that if enough cases had been done, some difference would have been detected.

Wavelet compression methods appear to perform better than JPEG, particularly for large-matrix images like radiographs.^{27,41} Using the dual pathology approach, reports of compression ratios of as high as 80:1 showing equal or better performance than originals has been reported.¹² Imaging modalities such as MR, CT,²² and nuclear medicine require much lower compression ratios, and the clear advantage of wavelet over JPEG is less clear.

IRREVERSIBLE COMPRESSION AND DICOM

As more studies evaluating compression techniques and their effects on diagnostic and aesthetic appearance are performed, wider acceptance of irreversible compression will occur. Adoption of a commercial standard will make supporting a irreversible compression method more appealing to vendors. Working Group 4 of the DICOM standards committee (image compression) has a liaison to the JPEG 2000 committee for the purpose of maximizing the usefulness of JPEG2000 for medical imagery, so that there will be a well-recognized and supported technology for compression of medical images.

There now is good evidence that irreversible compression can be used for medical image storage and transmission without compromising diagnostic value. However, it must be used in carefully managed, well-understood ways. An important step toward achieving this is to have standardized methods for compression. ACR-NEMA version 2.0 supports JPEG irreversible compression, and there has been some effort to extend Digital Imaging and Communications in Medicine (DICOM) to support wavelet compression. Working Group 4 (Compression) of the DICOM standards committee has representation on the JPEG2000 committee, and JPEG2000 has been approved as a DICOM standard, with ongoing work to adopt future JPEG2000 features into DICOM. Wavelet methods appear superior to JPEG for medical images, but the lack of a single standard wavelet method has hampered comparison of results and cross validation. Adoption of JPEG2000 as a DICOM standard should hasten the application of wavelet compression in medicine.

APPLICATION OF IRREVERSIBLE COMPRESSION TO MEDICAL PRACTICE

The most common application of irreversible compression in radiology is teleradiology. This application benefits tremendously because of the low bandwidth connections most homes have. Although technologies like cable modems and-digital subscriber lines (DSL) have increased bandwidth substantially, the need for compression seems to remain. We have found that over the course of 25 years, the size (in bytes) of a 'typical' CT examination (as well as MR over the past 20 years) closely parallels the increase in speed of wide-area connections (Fig 5). In clinical practice, this means that the typical emergency CT head for subarachnoid hemorrhage has become the nonenhanced CT followed by the contrast-enhanced CT angiogram or an MRI with diffusion and perfusion images.

Another application of compression is to reduce the storage and bandwidth requirements required to deliver images to clinicians.⁴² In this system, all electronic images are transmitted to a server where they are compressed and stored. Because of the high compression ratio, several months of images are available immediately

using a standard server. In addition, because images are decompressed at the desktop, the network bandwidth required to distribute these images is reduced substantially. Finally, images are stored in RAM and compressed, significantly reducing desktop computer resources demands. At the Mayo Clinic, this has reduced infrastructure costs by approximately \$250,000 per year.⁴³ In this application, image quality is satisfactory, and clinical acceptance has been overwhelmingly positive.^{44,45}

Medical-legal uncertainties are a significant hurdle to widespread use of irreversible compression for diagnosis. It can be argued that compression is not substantially different from any other step in the image creation and presentation chain. Just as radiologists routinely accept the trade-offs of using fewer excitations in an MR acquisition to reduce imaging time, and do not routinely double-process x-ray films, we also should be able to accept irreversible compression when it is shown that degradation of diagnostic accuracy does not occur.

The Food and Drug Administration finds that the use of irreversible compression is acceptable but must be noted when compressed images are displayed. They have now proposed a classification scheme for medical image management devices developed, which accounts for irreversible compression.⁴⁶ Within general guidelines, the technique and ratio are left to the radiologist's discretion, although the use of irreversible compression must be noted. Irreversible compression with a guarantee of diagnostic acceptability could allow PACS implementation with substantially fewer resources: the storage system for images would be substantially smaller, and the network bandwidth demands would be less if images were decompressed at the workstation.

CONCLUSIONS

Irreversible compression appears to be a very effective means of decreasing image file size to facilitate storage and transmission of radiologic images. There is increasing evidence that some forms of irreversible compression can be used with no measurable degradation in aesthetic or diagnostic value. It is increasingly necessary for radiologists to become conversant in compression techniques and their effects on

images. Using irreversible compression in everyday practice can reduce significantly the cost of delivering radiology services by reducing the information infrastructure required to deliver and store images.

REFERENCES

1. Karson TH, Chandra S, Morehead AJ, et al: JPEG compression of digital echocardiographic images: Impact on image quality. *J Am Soc Echocardiography* 8:306-318, 1995
2. Wallace GK: The JPEG still picture compression standard. *Comm of the ACM* 34:30-44, 1991
3. Gillespy T 3rd, Rowberg AH: Displaying radiologic images on personal computers: Image storage and compression — Part 2 [published erratum appears in *J Digit Imaging* 1994 2:60]. *J Digit Imaging* 7:1-12, 1994
4. Manduca A: Interactive wavelet-based 2-D and 3-D image compression. *Proc SPIE* 1897:307-318, 1993
5. Antonini M, Barlaud M, Mathieu P, et al: Image coding using wavelet transform. *IEEE Trans Image Proc* 1:205-220, 1992
6. Villasenor J, Belzer B, Liao J: Wavelet filter evaluation for image compression. *IEEE Trans Image Proc* 4:1053-1060, 1995
7. Shapiro JM: Embedded image coding using zerotrees of wavelet coefficients. *IEEE Trans Signal Proc* 41:3445-3462, 1993
8. Said A, Pearlman W: Image compression using the spatial-orientation tree, *in* *IEEE Intl Symp on Circuits and Systems*. Piscataway, NJ, IEEE Press, 1993, pp 279-282
9. Said A, Pearlman W: A new fast and efficient image codec based on set partitioning in hierarchical trees. *IEEE Trans Circuits and Systems for Video Technology* 6:243-250, 1996
10. Donoho DI, Johnstone IM: Ideal denoising in an orthonormal basis chosen from a library of bases. *Comptes Rendus Acad Sci.* 319:1317-1322, 1994
11. Erickson BJ, Manduca A, Persons KR, et al: Evaluation of irreversible compression of digitized PA chest radiographs. *J Digit Imaging* 10:97-102, 1997
12. Savcenko V, Erickson BJ, Palisson PM, et al: Detection of subtle abnormalities on chest radiographs after irreversible compression. *Radiology* 206:609-616, 1998
13. Slone R, Foos D, Whiting B, et al: Assessment of visually lossless irreversible image compression: Comparison of three methods by using an image-comparison workstation. *Radiology* 215:543-553, 2000
14. Persons KR, Palisson P, Manduca A, et al: An analytical look at the effects of compression on medical images. *J Digit Imaging* 10:60-66, 1997
15. Erickson B, Campeau N, Huston J, et al: Effects of irreversible compression on Neuro CT and MR Images. *American Society of NeuroRadiology*, Atlanta, GA, Apr 3-8, 2000
16. Girod B: What's wrong with mean-squared error? Watson A (ed): *Digital Images and Human Vision* (ed1). Cambridge, MA, MIT Press, 1993, pp 207-220
17. Barten P: Evaluation of the effect of noise on subjective image quality. *Proc SPIE* 1991 1453:2-15, 1991

18. Nill N, Bouzas B: Objective image quality measure derived from digital image power spectra. *Optical Engineering* 31:813-825, 1992
19. Good W, Lattner S, Maitz G: Evaluation of image compression using plausible "non visually weighted" image fidelity measures. *Proc SPIE* 1996 2707:301-309, 1996
20. Lattner S, Good W, Maitz G: Visually weighted assessment of image degradation resulting from image compression. *Proc SPIE* 1996 2707:507-518, 1996
21. Ishigaki T, Sakuma S, Ikeda M, et al: Clinical evaluation of irreversible image compression: Analysis of chest imaging with computed radiography. *Radiology* 175:739-743, 1990
22. Goldberg MA, Gazelle GS, Boland GW, et al: Focal hepatic lesions: Effect of three-dimensional wavelet compression on detection at CT. *Radiology* 202:159-165, 1997
23. Clunie DA, Mitchell PJ, Howieson J, et al: Detection of discrete white matter lesions after irreversible compression of MR images. *Am J Neuroradiol* 16:1435-1440, 1995
24. Good WF, Maitz GS, Gur D: Joint photographic experts group (JPEG) compatible data compression of mammograms. *J Digit Imaging* 7:123-132, 1994
25. Cosman PC, Gray RM, Olshen RA: Evaluating quality of compressed medical images: SNR, subjective rating, and diagnostic accuracy. *Proc IEEE* 82:919-932, 1994
26. Elion JL, Whiting JS: Clinical use of lossy image compression in digital angiography [editorial; comment]. *Am J Cardiol* 78:219-220, 1996
27. Goldberg MA, Pivovarov M, Mayo-Smith WW, et al: Application of wavelet compression to digitized radiographs. *AJR Am J Roentgenol* 163:463-468, 1994
28. Rebelo MS, Furuie SS, Munhoz AC, et al: Lossy compression in nuclear medicine images. *Proceedings of the Seventeenth Annual Symposium on Computer Applications in Medical Care*, New York, NY, 1993
29. Foos D, Muka E, Slone R, et al: JPEG 2000 compression of medical imagery. *The International Society For Optical Engineering (SPIE)*, San Diego, CA 2000
30. Savchenko V, Erickson B, Persons K, et al: An evaluation of JPEG and JPEG 2000 irreversible compression algorithms applied to neurologic computed tomography and magnetic resonance images. *J Digit Imaging* 13:183-185, 2000 (suppl 1)
31. Uchida K, Nakamura K, Watanabe H, et al: Clinical evaluation of irreversible data compression for computed radiography in excretory urography. *J Digit Imaging* 9:145-149, 1996
32. Wenzel A, Gotfredsen E, Borg E, et al: Impact of lossy image compression on accuracy of caries detection in digital images taken with a storage phosphor system. *Oral Surgery, Oral Medicine, Oral Pathology, Oral Radiology, & Endodontics* 81:351-355, 1996
33. Sayre J, Aberle DR, Boechat I, et al: Effect of data compression on diagnostic accuracy in digital hand and chest radiography. *Proc SPIE* 1653:232-240, 1992
34. Mori T, Nakata H: Irreversible data compression in chest imaging using computed radiography: An evaluation. *J Thorac Imaging* 9:23-30, 1994
35. MacMahon H, Doi K, Sanada S, et al: Data compression: Effect on diagnostic accuracy in digital chest radiography. *Radiology* 178:175-179, 1991
36. Breeuwer M, Heusdens R, Gunnewiek RK, et al: Data compression of x-ray cardio-angiographic image series. *Int J Cardiac Imaging* 11:179-186, 1995
37. Onnasch DG, Prause GP, Ploger A: Objective methods for optimizing JPEG compression of coronary angiographic images. *Int J Cardiac Imaging* 11:151-162, 1995
38. Yamamoto LG: Using JPEG image compression to facilitate telemedicine. *Am J Emerg Med* 13:55-57, 1995
39. Aberle DR, Gleeson F, Sayre JW, et al: The effect of irreversible image compression on diagnostic accuracy in thoracic imaging. *Invest Radiol* 28:398-403, 1993
40. Chan HP, Lo SC, Niklason LT, et al: Image compression in digital mammography: Effects on computerized detection of subtle microcalcifications. *Med Physics* 23:1325-1336, 1996
41. Saittech P, Ho BK, Panwar R, et al: Applying wavelet transforms with arithmetic coding to radiological image compression. *IEEE Eng in Medicine and Biology* 14:587-593, 1995
42. Erickson BJ, Ryan WJ, Gehring DG, et al: Image display for clinicians on medical record workstations. *J Digit Imaging* 10:38-40, 1997
43. Erickson B, Persons K, Hangiandreou N, et al: Requirements for an enterprise digital image archive. *J Digit Imaging* 14:72-82, 2001
44. Erickson B, Ryan W, Gehring D, et al: Clinician usage patterns of a desktop radiology information display application. *J Digit Imaging* 11:137-141, 1998
45. Eversman W, Pavlicek W, VZavalkovskiy B, et al: Performance and function of a desktop viewer at Mayo Clinic Scottsdale. *J Digit Imaging* 13:147-152, 2000 (suppl 1)
46. Food and Drug Administration: Federal Register 61, 1996