CHAPTER

12

IMAGE COMPRESSION

BRADLEY J. ERICKSON

Implementation of any picture archiving and communication system (PACS) requires a discussion of how to deal with the large quantities of data that must be transmitted and stored. Even before PACS, clever engineers devised encoding schemes that reduced the apparent size of images to reduce the demands placed on transmission and storage devices. Encoding images in order to reduce storage and transmission demands is called "image compression." Compression of medical images is controversial—some argue that it should not be used, because soon storage and networks will be cheap enough that compression is not necessary. This argument has been made from the early days of PACS and continues to be made. I suspect this issue will continue to be controversial until humans no longer view images.

Those faced with the need to implement systems in the present point to the benefits that can be achieved. Sometimes the argument for using compression is purely financial: using compression tips the scale from a losing proposition to a winning one. It can allow one to use less expensive networking technologies and requires less (potentially much less) storage. These can be major components of the capital acquisition as well as operating costs.

In other cases, compression can make the difference between a project that is feasible and one that cannot work. When a specific turnaround time is required to achieve satisfactory service and the network capacity is limited, compression can be the only solution. This is most often the case when remote facilities are an integral part of the PACS equation.

Whether cost savings or turnaround time is the issue, the counter-argument is frequently made that if one delays the project, the advance of technology can obviate the need for compression. This is an appealing argument, for Moore's law seems to apply not only to processor speed but also to network and storage devices. Unfortunately, this oversimplifies the situation. The primary reason this argument fails is that the demands on the imaging system are not static. The standards for PACS performance continue to rise in step with Moore's law. That is because the same advances in processing power that make a PACS cheaper are also making imaging devices that produce higher image data volumes. A few years back, we reported that over a long period of time, the daily image data volume for computed tomography (CT) and magnetic resonance (MR) scanners was paralleling the increases in network speed and storage density. Figure 12.1 shows that the relationship has held true for 30 years. The conclusion is that one should not delay a project in expectation that technology advances will eliminate the need for compression. *If a PACS implementation is not viable without compression today, it likely will not be viable in the future. Similarly, if compression is*

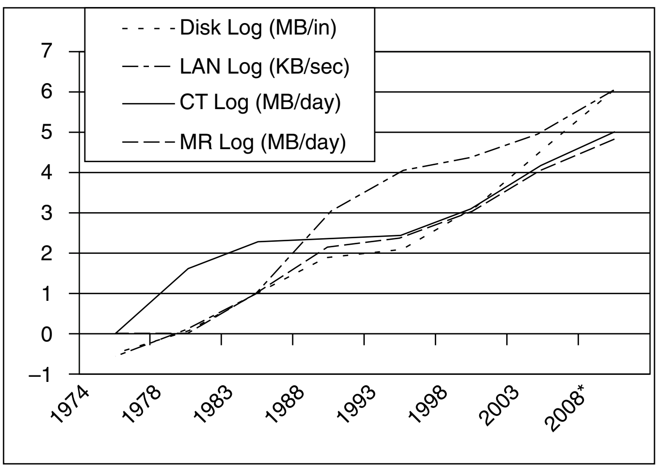

FIGURE 12.1

Graph of network speed, storage density, and daily CT and MR data volumes from 1974 to 2004. * = projected.

necessary at the start of a project, one should plan that it will continue to be necessary for its lifetime.

IMAGE COMPRESSION TECHNOLOGIES

There are two main categories of image compression: reversible, or lossless, and irreversible, or lossy. Reversible compression will exactly reproduce the original image after compression and decompression. There is little controversy surrounding the use of lossless compression—the major question is whether the computational effort is worth the modest gain. For most medical images, the gain is on the order of 2 : 1 to 3 : 1. A 2 : 1 ratio means that the storage or network demands are 50% of what the original image would require; 3 : 1 would require just 33%. Many systems implement this type of image compression because the computation cost is small compared to the cost savings in storage and networking. In some cases, this compression occurs only when the image is saved to the long-term archive, saving only storage costs. Some storage devices actually have hardware that performs this type of compression.

There are 2 main ways that lossless compression algorithms work. The first type is the group that utilizes repetition within an image to compress the image. For instance, if there were several consecutive pixels with the same value, it is faster to say "5 pixels of value 312" than to say "312 312 312 312 312." Other than some outer regions of images (e.g., the pixels that fill the circle of a CT to make a square) or three-dimensional (3-D) renderings, exact repetition of pixel values is rather rare for medical images. However, an adaptation is to use the prior value as a starting point and just send the difference. Other than where there are strong edges, small numbers (which take less space) added to the prior value will correctly represent the pixel value.

An alternative to this strategy is to look for patterns within the image. If one were to look through this text, there would be a fairly high frequency of "the" and "to," and one would find that "q" was always followed by "u." If a single-letter code were substituted for these patterns, the storage space required for the text could be significantly reduced. This is a dictionary-style compression method and is commonly used for highly structured information, such as text. It is used in most of the "zip" programs that are available to compress files.

The lines that follow this paragraph show examples of lossless compression methods. The original pixel values are shown in [a]. A run-length encoding of the pixels could be that shown in [b]. Note that there are more values sent in line [b] than line [a], though the value encoding how many of a kind may not take as much space as the value. Line [c] shows a differential encoding of the pixels. The number of values is the same, but if the change values are small, they will take less space to transmit. The next line, [d], has two parts: a dictionary and the data. The data values are used to look up the actual pixel values for the output. The first entry in the dictionary is the most common pattern and uses the fewest bits possible, with less frequently used patterns consuming more bits. Note that in all of these examples, the compressed data is larger than the original. This is because in such a small dataset there is little redundancy to take advantage of and because different data sizes (bits) are not represented.

```
[a] 1   4   4   4   5   6   7   8   8   4   5   4   5   1   4
[b] 1   1)(3   4)(1   5)(1   6)(1   7)(2   8)(1   4)
    (1   5)  (1   4)(1   5)(1   1)  (1   4)
[c] 1   3   0   0   1   1   1   1   0  -4   1  -1   1  -4   3
[d] dictionary:
    (4   5)  (4   4)  (1)  (6)  (7)  (8   8)  (4)
    data:
    3   2   1   4   5   6   1   1   3   7
```

While there is some structure to medical images, the gain achieved by compression of images is not as great as that achieved by compression of text. Compared with the repetition strategy described above, pattern-based compression typically will achieve higher compression ratios but requires more central processing unit (CPU) power and memory to compress and decompress. However, one can devise methods in which most of the CPU power is used for the compression task (which is done just once) and less for decompression (which is done every time the image is viewed). Unlike lossless compression, where the compression ratio is the same for a given image, the implementation of lossy compression requires decisions about how much compression to apply. Depending on the algorithm, one may specify a "quality factor," which roughly correlates with perceived image quality after compression, or the algorithm may require a specific compression ratio. It would seem that the quality factor approach is more desirable, as each image would then be optimally compressed. In practice, however, the correlation with perceived image quality is variable. This will be discussed in greater depth later in this chapter.

Lossy compression begins with a very different step from lossless compression that is not very intuitive. That first step is a transformation, and the type of transformation is often used to characterize the type of compression. An early popular transformation was the discrete cosine transform (DCT). In some cases, the entire image had this transform applied. However, this is computationally intensive, particularly for the computers of the 1980s, when the technique gained popularity. Therefore, an industry group known as the Joint Photographic Experts Group (JPEG) created a processing algorithm for images that would break the image into 8×8 pixel blocks and apply the DCT to each block. These smaller blocks were more amenable to computers of the day and permitted dedicated hardware to be built for this task. In either case, the DCT would create an image (with the same matrix size as the original image/block) in which the low-frequency components were concentrated in the lower left corner, with the highest-frequency elements in the upper right. The values in this transformed image matrix are sometimes referred to as coefficients because they are coefficients for a mathematical formula that can re-create the image. More recently, the various forms of the wavelet transformation have been applied for image compression. Such wavelet transforms produce multiple images at different scales, where the next-higher-resolution image has the higher-frequency changes from its lower-resolution predecessor. Figure 12.2 shows an example of a 4-level wavelet transformed image.

After the transformation of choice has been performed, the image has been altered, but there has been no loss of information—one could invert the transformation and exactly reproduce the original image. The next step after transformation is quantization, which is the step of deciding which parts of the transformed image are important and which are not. One could select a low-frequency rectangle and compute the inverse transform of it to re-create an image that would be somewhat blurrier than the original. How blurry the resulting image is depends on how effectively the transformation concentrated information into the selected parts of the transformed image.

There are 2 primary ways one can save storage/transmission space: by assuming that coefficients not sent are 0 (zero-filling) or by reducing the accuracy of the coefficients that are sent (e.g., converting from 80-bit floating point values to 20-bit scaled integers). In some schemes, including JPEG, one can also preferentially save accuracy on those coefficients that represent frequencies that the human eye is sensitive to, while saving bits on those components to which the eye is not sensitive. The multiresolution nature of wavelets also permits more sophisticated selection of parts of the image than just a rectangle. Some algorithms, such as Set Partitioning in Hierarchical Trees (SPIHT), efficiently represent areas of images that have much activity while saving bits in uninteresting areas. This is complementary to

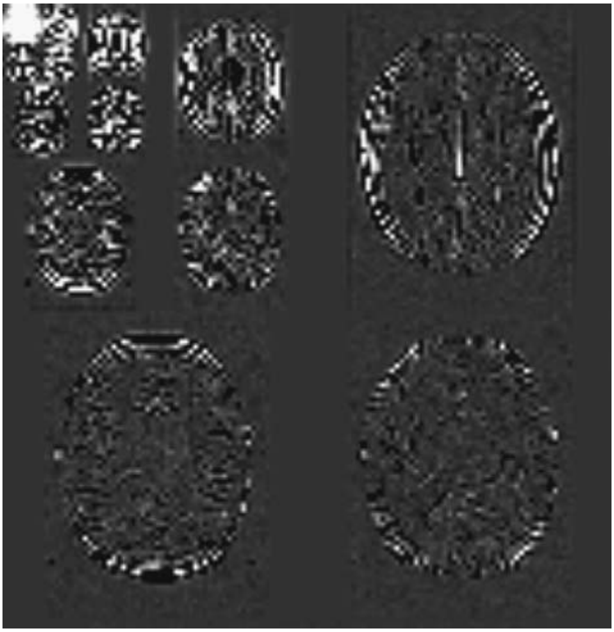

FIGURE 12.2

Example of a wavelet transformed image. The lowest-frequency components are in the upper left corner. The original image can be exactly reconstructed from this image.

preferring spatial frequencies to which the eye is sensitive, as described above.

In some cases, it is also desirable to sequence the coefficients so that those that are most important to recreating a faithful image are first in the stream of data, rather than ordered by an X, Y position. This can allow the user to get a quick view while the image continues to improve as more information is sent. This is called progressive streaming (referring to the order of data transmission) or successive refinement (referring to the actual displaying of improved images).

After the coefficients have been quantized, there is often still some redundancy, and so the third step in lossy compression is to use one of the lossless techniques described above (typically a dictionary-type method) to further compress the quantized coefficients.

Unlike lossless methods, in which the compression ratio is fully determined by the algorithm selected, lossy methods allow the user to select either the compression ratio or a quality factor. In either case, the user's choice affects how much the coefficients are quantized, meaning that a judgment is made about acceptable compression level. Unfortunately, in most cases it is not feasible for the user to interactively select the appropriate compression for each and every image, and so general rules for compression must be used.

APPLICATION TO MEDICAL IMAGING

The description of image compression provided above is not unique to medical images—these image compression algorithms are used for images on the Internet or from your digital camera. However, some unique aspects of medical images deserve comment. The first is that most medical images are grayscale, with more gray levels than most commercial/consumer images. Radiographic images are 10 to 12 bits, while CT and MR images are 12 to 16 bits. Since computers prefer to deal with pixels that are multiples of 8 bits, most medical images are stored using 16 bits per pixel. One exception is ultrasound. In some cases (particularly Doppler), ultrasound images use 24-bit color, which is like that used in commercial/consumer imaging. Grayscale ultrasound images are variable and may be either 8-bit or 16-bit. The significance is that JPEG was created without attention to medical imaging, and the extensions to JPEG to allow it to support 12 or 16 bits are not always compatible.

Another difference is that CT and MR images are often from spatially adjacent locations, meaning there is significant redundancy between 2 or more images. This is not considered in any of the (two-dimensional [2-D]) algorithms described above. However, an update of the JPEG standard (referred to as JPEG2000) does have components that support the notion of a 3-D stack of slices. While the JPEG2000 standard is still advancing, the basic 2-D methods have been ratified, and there are implementations that have been tested on medical images. Some of those results will be discussed later in this chapter. Because compression has been considered important to the success of PACS by many of its leaders, there was a concerted effort to participate in the JPEG2000 standard in order to make it better support medical images—in particular, by addressing the 16-bit issue. The Digital Imaging and Communications in Medicine (DICOM) standard has included JPEG-based compression in its transfer syntaxes for some time. The approved part of JPEG2000, which focuses on 2-D image compression using a wavelet transformation, has already been adopted as a DICOM standard (Supplement 61, November 2001). It includes support for 16-bit images, progressive encoding, and region-of-interest functions (improved quality in specific regions of the image). It is expected that as other parts of JPEG2000 are ratified, such as the volumetric image functions (JP3D or JPEG2000 Part 10) and Motion JPEG for endoscopy, DICOM will add them to its standard.

EVALUATION IN MEDICAL IMAGING

IS IT WORTH IT?

There is controversy over whether there is any acceptable level of lossy compression for medical images or if there should be no allowance for alteration of pixel data. I believe that lossy compression can be acceptable and justified. Many major medical centers have come to the same conclusion. There are several bases for this conclusion.

First, as with nearly all parts of medical practice, one must exercise judgment about the trade-off of benefits and risks. Higher-quality radiographs can be obtained if higher-exposure images are obtained. Whether this would improve diagnosis or simply reduce quantum noise is the judgment we make when an exposure/dosage decision is made. In the case of MR and CT, one could reconstruct many more images than is typically required using interpolation. In the case of MR, one could use many more signal averages to reduce noise. But in order to conserve resources (reduce MR examination time or reduce total images per examination) we make judgments about how much is enough to be diagnostic. The same is true with image compression. As long as there is a thorough understanding of the trade-offs, and confidence that diagnostic accuracy is not significantly impacted, compression is no more "dangerous" than other decisions in medical imaging.

Second, the pixel data altered by compression has nearly always been altered even before it is compressed. CT uses reconstruction filters that may blur to reduce noise. Computed radiographs have lookup tables applied to emphasize certain regions of the histogram, while suppressing others, and may also have spatial filtering. These are also carefully considered alterations of pixel data in order to improve radiologist efficiency. Perhaps the greatest loss of data occurs in fluoroscopy, in which only a few spot images are taken from several minutes of an examination.

Third, it has been shown that at modest levels, lossy compression actually improves diagnostic performance because it preferentially loses the noise. This has been shown to improve detection of breast cancers as well as lung pathologies. From this perspective, compression could be viewed as a valuable processing step. This also raises issues about implementation. If lossy compression makes a lesion more conspicuous, compressing an image only after interpretation could cause one to miss a lesion that is more apparent with compression. If the lossy compressed version is the only version archived, it also could lead to greater difficulties if there should be any medico-legal issues about lesion perception/diagnosis. For this reason, when our department had the opportunity to implement lossy compression prior to diagnostic review, we did so. In that case, no one ever saw the original images—the same lossy compressed image is interpreted by the radiologist, distributed to the referring physician, and archived for any later review. Of course, this was done only after a large-scale review of the compression algorithm that was implemented.

In these days when containment of medical costs is receiving much attention, we should not be reluctant to make a carefully considered, thoroughly studied decision to reduce costs by implementing lossy compression.

SETTING THE STANDARD

If one agrees that using lossy compression can be an acceptable decision, the next question is how to determine what is good enough. As noted in the technical description of lossy algorithms, nearly all allow the user to select at least the range of compression that is achieved. That is really the challenge of compression: how to maximize the compression ratio without degrading the image. At least 4 methodologies have been used to answer this question.

The first method is to use human visual ratings. In this case, one simply asks radiologists either: (1) Are the compressed images good enough? or (2) Is there any difference between compressed and uncompressed images? The first question is a bit more difficult to justify, as it involves a judgment about what is required for diagnosis. In many cases, the abnormalities are so obvious that even marked image degradation could still permit a diagnosis. This is usually how x-ray exposure is selected. An alternative that seems more acceptable is to see whether there is any perceptible difference between compressed and uncompressed images. If one cannot tell the difference, it is unlikely to make a diagnostic difference. This type of study is easy to execute, since one can select randomly from all images and then present pairs of images (the original and the compressed/decompressed) to a blinded rater. If the rater prefers the compressed images at the same rate (or higher rate due to de-noising effect), then the compression ratio is acceptable.

The second method seeks to prove that compressed images are diagnostically equivalent to uncompressed. To do this, one must have proof of what the diagnosis is for each case. Furthermore, one should really have a range of diseases that would be expected to be diagnosed with that type of image. Ideally, the gold standard for knowing that disease is present should be established by an independent method, since we know that compression can improve performance. If an independent standard were not used, all the correct diagnoses made possible by improved signal-to-noise ratio would be counted as misses. And finally, you must have enough of each type of disease to prove a difference. Note that failing to find a difference is not the same as proving equivalence. The latter is much more demanding statistically, while failing to find a difference is a certainty if you have only a few cases.

One example of this methodology is a study we did with chest radiographs, in which we selected 2 common pathologies: nodules and fibrosis. We selected only 2 diseases to keep the receiver operating characteristic (ROC) design simple. We also selected them because they had very different properties—uncalcified nodules are low-frequency findings, while fibrosis is a high-frequency finding. Chest CT was used as the gold standard. Mammography can employ the same technique—masses versus microcalcifications, with biopsy or clinical follow-up being the gold standard.

The third method is to use computer-aided diagnosis (CAD) output as a detector for diagnostic degradation. This has the advantage of allowing many cases to be assessed with less effort. In addition, the variability of a human observer is removed. The major concern is that the features important to a CAD algorithm may not be the same for a human. Furthermore, CAD algorithms are generally focused on detecting a single disease (e.g., cancer). In cases like mammography, where cancer is nearly the entire disease range, this is not a problem, but for nearly all other modalities, preservation of CAD performance ensures only that features important to that CAD algorithm are preserved, not necessarily all the features necessary for a radiologic interpretation.

The fourth evaluation method is also computer-based but uses a model of the human visual system to predict when lossy compression has produced alterations that would be visible to a human. Like method 3, this method has the advantage of allowing many cases to be evaluated with little human effort. It has the additional advantage over the CAD method in that it should be valid for any type of image or disease. However, this method is complex to implement—one must have a valid model of the human visual system, which, in turn, requires knowledge of the output device. The computations for this are also nontrivial. Nevertheless, there are now some reports that are using this method.

COMPRESSION APPLICATIONS IN MEDICAL IMAGING

Clearly, the goal of medical image compression is to reduce the amount of data that must be stored or transmitted. Depending on the application, either lossy or lossless compression may be the best choice. Rather than review applications using this division, it is probably more appropriate to evaluate compression for each modality. This is because the compressibility of images is best correlated with modality.

MAMMOGRAPHY

Mammography is an area that has received much attention from the compression industry. That is because the data sets are large (40 megabytes [MB] per image × 4 images per screening exam), there is a large exam volume because mammography is used for screening a large at-risk population, and expertise is often not distributed to areas that acquire the images. This combination of factors means that moving the images from the acquisition location to a distant location where there is expertise in interpreting mammograms is often necessary.

One factor that has made analysis of mammogram compression easier is that there is really only one disease process of interest: breast cancer. However, breast cancer can have any of 3 types of appearance. The first is that of a small area of tiny calcifications, also known as microcalcifications. The second appearance is of a mass, which will often have irregular borders. The third is distortion of the architecture of the tissue within the breast. Because these 3 appearances can be fairly easily characterized, the analysis of compression effects is simpler than for other image types.

Early mammography compression studies used digitized films and either JPEG or a wavelet variant. Because early algorithms tended to lose high-frequency information first, most attention was paid to the effect on microcalcifications. Studies from the 1990s showed that JPEG compression could be applied at levels up to 15 : 1 without producing perceptible changes in the images, as assessed by image-processing experts. One should note that this study did not use radiologists to evaluate the images. This is often a problem—to do a good study, one must have many cases with subtle disease with an independent gold standard, and these many cases should be evaluated by radiologists. Just compiling a large set of subtle cases with proof of diagnosis is difficult. Getting several experts to review these cases is more difficult. And since the cases selected are subtle, there will be significant variability in the ratings—variability that will obscure any subtle difference caused by compression.

For these reasons, some have begun using computer-aided diagnosis (CAD) as a measure of acceptable compression. This is an elegant way to avoid the variability inherent in human observers, and to evaluate the large number of cases required to detect small differences.

Early studies, of necessity, used digitized film mammograms. Digital detectors are beginning to appear on the market. It is possible that the image properties are different, and so it is necessary to perform separate studies on these images to determine the correct compression ratio. But at the same time, compression algorithms, such as JPEG2000, are advancing. JPEG2000 is receiving much attention because it has been adopted into the DICOM standard and because it holds promise for better compression performance than standard JPEG. One recent study using an alternative forced-choice method found that digital mammograms compressed with JPEG2000 at ratios up to 20:1 were indistinguishable from the originals, which is similar to results for digitized films.

COMPUTED RADIOGRAPHY/DIGITAL RADIOGRAPHY

Compression of radiographic images also has a long history, beginning with study of digitized radiographs but now focusing on computed radiography (CR) and digital radiography (DR) images. It is encouraging that the results seem similar: digitized radiographs seem to be about as compressible as CR and DR. For chest radiographs, ratios in the range of 20:1 seem to produce no visible or diagnostic degradation. Slone found that with very close inspection/magnification, images compressed at 10:1 using JPEG are still indistinguishable from originals, and at normal viewing conditions, 20:1 is equivalent. Compression at ratios of up to 32:1 (either JPEG or wavelet) did not degrade detection of simulated nodules on chest phantom images. We found that for either nodule diagnosis or interstitial disease detection, wavelet compression of digitized chest X-rays (CXRs) at up to 40:1 was not significantly different for nodules and for interstitial disease, and performance at 10:1 showed a trend to be superior to original images. There are fewer studies of musculoskeletal radiographs, but those that exist also show that compression in the range of 20:1 does not alter visual appearance.

COMPUTED TOMOGRAPHY AND MAGNETIC RESONANCE IMAGES

Images with a smaller matrix, such as CT and MRI, might be expected to be less compressible because there is less potential for redundancy; this seems to be the case. Whereas radiography seems to tolerate ratios in the 20:1 range, acceptable compression ratios for CT and MR seem to be closer to 10:1. There is also very little difference between compression methods, likely because there is less redundancy for full-frame methods such as JPEG2000 and wavelet to leverage. In a study of wavelet compression, ratios of 8:1 did not affect accuracy of diagnosis of acute appendicitis, but 16:1 and 24:1 did show decreased sensitivity. Another study showed no change in nodule detection on low-dose chest CT at 10:1 utilizing wavelet compression; 10:1 was also found equivalent for both JPEG and wavelet for detection rate for lung cancers on low-dose chest CT. Brain MRIs compressed using a wavelet algorithm at up to 20:1 showed no difference in ROC value for variety of lesions for a 512×256 matrix and 10:1 for a 256×256 matrix. This reflects rather nicely that the higher matrix had little additional information. JPEG compression of head CTs at up to 20:1 did not degrade ROC performance for detecting infarction, though infarction is a low-frequency finding that may be more compression tolerant.

While it would be nice to have a single ratio for a modality, this is not the case. Conventional chest CTs compressed at 6:1 (JPEG) were considered acceptable, but only 4:1 was acceptable for thin section (2 mm) CT. This is an interesting finding that despite the fact that compression seems to discard noise preferentially, noisy images are less compressible. It demonstrates the importance of redundancy and texture versus information in visual appearance. This is an important problem because of the rapid expansion of multidetector CT. It is fairly simple to create very thin images as well as thicker images for a given body part. Studies comprising thousands of images could become routine. The fact that the greater noise in these thin sections dominates the redundancy is significant—not only are there more slices to store and transmit, but they are also less compressible. Since potentially multiple datasets are derived from the same projection data, some have proposed that it may be more efficient to compress the projection images than to compress the reconstructed sections.

FLUOROSCOPY

The largest systematic study of fluoroscopic compression, and probably of any medical imaging modality, is for cardiac angiograms. This large, multisite trial, which focused on the visibility and appearance of coronary artery stenosis, showed that JPEG compression at 6:1 was equivalent to the original images, with some degradation of quality and performance at 10:1. It is reasonable to expect that other fluoroscopic images will exhibit similar compressibility.

NUCLEAR MEDICINE

Nuclear medicine images also comprise a wide range of image types. Some are very low resolution (64 × 64), while others are similar to CT. We have found the small matrix images to be rather incompressible, and in those cases, we simply use lossless compression. It is similar to the text example above—there are so few samples to work with that finding redundancy is difficult. For 256 and up matrix images, we have found compression ratios similar to CT and MRI (10:1) to be acceptable.

ULTRASOUND (STATIC AND VIDEO)

Diagnostic ultrasound actually produces images of several types: static grayscale images that may be captured video with 256 gray levels, static grayscale images with 1000 to 4000 gray levels, static color images with 24 bits of color, and real-time or video signals. Despite this wide variety of image types, all seem to allow compression on the order of 10:1, and this also seems to transcend compression algorithm. Some have also demonstrated no alteration in automated intimal wall thickness from ultrasound images after compression.

THE FUTURE

Information theory can measure how much unique information is present in an image. That should be the upper limit on lossless image compression ratios. However, there are (fortunately) many clever people who find that those limits can be broken by using additional information about the images to cheat the limits. One familiar example from CT is that the reconstruction typically is circular, while the image is square. Rather than encoding the information outside the circle, one could (and some scanner manufacturers did) never send any information for pixels outside the circle, thus improving the compression ratio. This is obviously an extreme case, but other information can be used to effectively reduce the data size without altering pixels of interest. For instance, CT data is viewed at narrow window widths over only a limited range of window levels. By altering the histogram, one can improve compression ratios without producing any perceptible image degradation *as normally viewed*.

THREE-DIMENSIONAL (3-D) IMAGING

Three-dimensional imaging modalities, such as helical CT, MR, and ultrasound (US), are those that are experiencing the greatest growth in data volumes. Since they can also produce 3-D data, improved compression rates may be achievable by utilizing methods that take advantage of the coherence that is through-plane, as well as in-plane. While early results with 3-D compression were disappointing, more recent attempts with thinner section data show more promise. Figure 12.3 shows a comparison of 2-D versus a videolike 3-D compressor versus a true 3-D spatial compressor. It is likely that this improved performance for 3-D in the later study is because the thinner sections possible with multidetector computed tomography (MDCT) allow more coherence of data between slices. However, the results continue to fall below expectations, likely because these thinner slices also have more noise, which decreases compressibility.

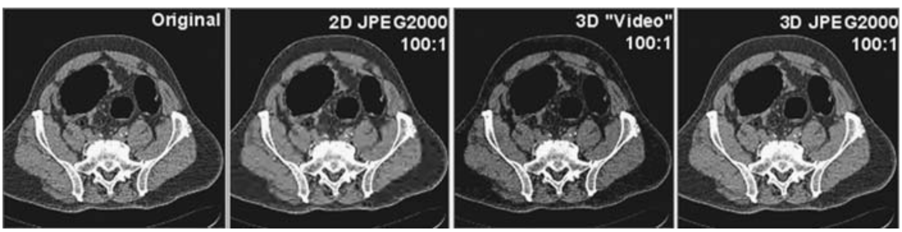

FIGURE 12.3

Comparison of JPEG2000 compression methods. The original image is the left-most panel. The next image is after 2-D JPEG2000 compression has been applied. The third uses a videolike compressor, and the fourth is 3-D JPEG2000 compression. All compressed images have a final compression ratio of 100:1. Note the marked degradation of the 2-D compressed image and the near-perfect 3-D compressed image. (Images courtesy of Michael Marcellin, University of Arizona.)

RAW DATA COMPRESSION

For modalities like CT and MR, the images that are reconstructed are only a small subset of the possible images that could be produced. Given this flexibility and the different personal preferences in how to view these increasingly complex image sets, it may be most effective simply to compress and store the raw data and the reconstruction parameters than to compress and store all the image products derived from the raw data. In the case of CT, this would mean compressing the projection data, while in MRI, it would mean compressing the k-space data. While there was some early investigation of these methods, recent trends may revive these techniques. While these examples are for CT and MRI, it is possible and even likely that multiple-image products will be created from a single source/raw dataset. For that reason, being able to compress the raw data, and describe (compactly) how the derivatives were made, will likely be more effective than storing all the derivatives.

HUMAN VISUAL SYSTEM–BASED COMPRESSION

A crucial element in the decision to use lossy compression is that there are no perceptible changes in the image, or that the changes are so small that they are of no diagnostic significance. Several studies have been done using the latter criteria, but they are somewhat less satisfying because there is always the nagging doubt of whether a particularly challenging case just might be compromised by compression. So-called visually lossless compression is more appealing—if the radiologist cannot see the difference, it is hard to imagine that a difference in diagnostic performance could exist. Operating on that assumption, there are some efforts under way to use models of the human visual system (HVS) to determine the threshold at which compression-related image alterations become perceptible. This means that each image is examined and its optimal compression ratio is computed and applied. This is probably more valid than general modality rules for acceptable compression ratios. Anyone who has applied lossy compression knows that there are always exceptional images that are very incompressible, for which general rules do not apply. These exceptional images would not be overcompressed with an HVS-based compression system. While an HVS system would be more expensive because of the need to compute the optimal ratio for each image, it might also save money by allowing a higher average compression ratio. Today, conservative ratios are applied to make sure most images are good enough. With HVS, *every* image would be good enough, and those that are very compressible would not be *under*compressed. Work on HVS-based compression is still preliminary, but such a system is a reasonable possibility and might address concerns that exist within the imaging community.

This chapter introduced compression by pointing out that it was controversial. But once computer algorithms are used to interpret images, there will be a quantitative measure of required image characteristics that will permit precise optimization of compression. This is directly analogous to modeling the HVS to determine acceptable compression. The difference is that HVS models are based on estimates of humans and specific viewing conditions—these may not match actual viewers and viewing conditions. With CAD, such variability is not present, enabling more confident statements that compression is not affecting image interpretation. Some may still argue that at some point, computers and networks will be inexpensive enough to render compression needless. The fact that this argument has been made for decades, and since Moore's law specifies an 18-month doubling time, I doubt this will be true in the foreseeable future. Indeed, the forces driving computer technology also drive medical imaging devices. Compression is a technology that we will have to grapple with for the foreseeable future. Much like x-rays themselves, compression can be applied to advantage as long as those employing it understand its strengths and weaknesses.

REFERENCES

Bae KT, Whiting BR. CT data storage reduction by means of compressing projection data instead of images: feasibility study. *Radiology*. 2001;219(3):850–855.

Brennecke R, et al. American College of Cardiology/European Society of Cardiology international study of angiographic data compression phase III. Measurement of image quality differences at varying levels of data compression. *Eur Heart J*. 2000;21(8):687–696.

Chan H, et al. Image compression in digital mammography: effects on computerized detection of subtle microcalcifications. *Med Phys*. 1996;23(8):1325–1336.

Good W, Maitz G, Gur D. Joint photographic experts group (JPEG) compatible data compression of mammograms. *J Digit Imaging*. 1994;7(3):123–132.

Erickson BJ, et al. Requirements for an enterprise digital image archive. *J Digit Imaging*. 2001;14(2):72–82.

Hangiandreou NJ, et al. The effects of irreversible JPEG compression on an automated algorithm for measuring carotid artery intima-media thickness from ultrasound images. *J Digit Imaging*. 2002;15(suppl 1):258–260.

Huang H. Progress in image processing technology related to radiological sciences: a five-year review. *Comput Methods Programs Biomed*. 1987;25(2):143–156.

Kerensky RA, et al. American College of Cardiology/European Society of Cardiology international study of angiographic data compression phase I. The effects of lossy data compression on recognition of diagnostic features in digital coronary angiography. *Eur Heart J.* 2000;21(8):668–678.

Ko JP, et al. Wavelet compression of low-dose chest CT data: effect on lung nodule detection. *Radiology.* 2003;228(1):70–75.

Kotter E, et al. Evaluation of lossy data compression of chest X-rays: a receiver operating characteristic study. *Invest Radiol.* 2003;38(5):243–249.

Li F, et al. Effects of JPEG and wavelet compression of spiral low-dose CT images on detection of small lung cancers. *Acta Radiol.* 2001;42(2):156–160.

Manduca A, et al. Histogram transformation for improved compression of CT images. In *Medical Imaging 1997*. Newport Beach, CA: SPIE; 1997.

Manduca A, et al. 3-D compression of medical images with Set Partitioning in Hierarchical Trees. In *RSNA*. Chicago: RSNA; 1997.

Megibow AJ, et al. Computed tomography diagnosis utilizing compressed image data: an ROC analysis using acute appendicitis as a model. *J Digit Imaging*. 2002;15(2):84–90.

Ohgiya Y, et al. Acute cerebral infarction: effect of JPEG compression on detection at CT. *Radiology*. 2003;227(1):124–127.

Persons KR, et al. Evaluation of irreversible JPEG compression for a clinical ultrasound practice. *J Digit Imaging.* 2002;15(1):15–21.

Ritenour ER. Lossy compression should not be used in certain imaging applications such as chest radiography. For the proposition. *Med Phys.* 1999;26(9):1773–1774.

Said A, Pearlman W. A new fast and efficient codec based on set partitioning in hierarchical trees. *IEEE Trans Trans Circuits and Systems for Video Technology.* 1996; 6:243–250.

Savcenko V, et al. Detection of subtle abnormalities on chest radiographs after irreversible compression. *Radiology*. 1998;206(3):609–616.

Siddiqui K, et al. Improved image compression at various slice thicknesses for multislice CT using 3-D JPEG2000 (part 2) in comparison with conventional 2-D compression. In *Soc Computer Applications in Radiology*. Vancouver, BC: SCAR; 2004.

Slone RM, Muka E, Pilgram TK. Irreversible JPEG compression of digital chest radiographs for primary interpretation: assessment of visually lossless threshold. *Radiology*. 2003;228(2):425–429.

Sung MM, et al. Clinical evaluation of compression ratios using JPEG2000 on computed radiography chest images. *J Digit Imaging*. 2002;15(2):78–83.

Suryanarayanan S, et al. A perceptual evaluation of JPEG 2000 image compression for digital mammography: contrast-detail characteristics. *J Digit Imaging*. 2004; 17(1):64–70.

Tuinenburg JC, et al. American College of Cardiology/European Society of Cardiology international study of angiographic data compression phase II. The effects of varying JPEG data compression levels on the quantitative assessment of the degree of stenosis in digital coronary angiography. *Eur Heart J.* 2000;21(8):679–686.

Yamamoto S, et al. Evaluation of compressed lung CT image quality using quantitative analysis. *Radiat Med.* 2001;19(6):321–329.

Zhang Y, Pham B, Eckstein M. Automated optimization of JPEG 2000 encoder options based on model observer performance for detecting variable signals in X-ray coronary angiograms. *IEEE Trans Med Imaging.* 2004;23(4):459–474.

Zheng B, et al. Applying computer-assisted detection schemes to digitized mammograms after JPEG data compression: an assessment. *Acad Radiol.* 2000;7(8):595–602.