

Joint Photographic Experts Group (JPEG) Compatible Data Compression of Mammograms

Walter F. Good, Glenn S. Maitz, and David Gur

We have developed a Joint Photographic Experts Group (JPEG) compatible image compression scheme tailored to the compression of digitized mammographic images. This includes a preprocessing step that segments the tissue area from the background, replaces the background pixels with a constant value, and applies a noise-removal filter to the tissue area. The process was tested by performing a just-noticeable difference (JND) study to determine the relationship between compression ratio and a reader's ability to discriminate between compressed and noncompressed versions of digitized mammograms. We found that at compression ratios of 15:1 and below, image-processing experts are unable to detect a difference, whereas at ratios of 60:1 and above they can identify the compressed image nearly 100% of the time. The performance of less specialized viewers was significantly lower because these viewers seemed to have difficulty in differentiating between artifact and real information at the lower and middle compression ratios. This preliminary study suggests that digitized mammograms are very amenable to compression by techniques compatible with the JPEG standard. However, this study was not designed to address the efficacy of image compression process for mammography, but is a necessary first step in optimizing the compression in anticipation of more elaborate reader performance (ROC) studies.

Copyright © 1994 by W.B. Saunders Company

KEY WORDS: Data compression, mammography, Joint Photographic Experts Group (JPEG), image processing, picture archiving and communication system.

IN RECENT YEARS, there has been a rapid increase in the use of mammography and a corresponding increase in the complexity of managing mammographic studies. Many believe that to accommodate the large number of mammographic studies now being performed, these images must be managed in a digital format. The main difficulty with this relates to the large volume of digital data needed to represent each image. High-resolution digital mammograms are likely to contain 10 to 40 Mbyte each. Digital-archiving and data-transmission systems quickly become strained when called upon to deal with a large number of these images. The application of image-compression techniques has the potential to greatly reduce this problem.

Because the problems of dealing with digital image data are universal, the Joint Photographic Experts Group (JPEG) has recently

defined a general-purpose image-compression standard.¹⁻³ Because the JPEG standard does not specifically address mammographic images, our aim in this study was to develop a compression technique for digitized mammograms which is compatible with the standard, but which is optimized for mammographic images, and then to determine the level of compression that can be achieved before the effects of compression become detectable by observers.

The JPEG algorithm is a block quantization version of the cosine transform in which images are divided into 8- × 8-pixel blocks and the cosine transforms applied to each block individually. The frequency coefficients calculated from the cosine transform are quantized by dividing by values from an array of quantization factors and rounding the quotients to integral values. The resulting array of quotients is aligned in a zigzag order and then encoded by either Huffman coding or by adaptive arithmetic coding. In this method, the quantization factors that are applied to the frequency coefficients determine the compression ratio as well as the kind of information lost or artifact introduced.

The main challenge in applying JPEG standard compression relates to the need to tailor the technique to the characteristics of the particular kind of images being compressed. There are two main areas where there is flexibility in applying the standard—in the quantization of the frequency coefficients and in preprocessing the image before the application of the JPEG algorithm.

In the past, most attention has focused on the

From the Department of Radiology, University of Pittsburgh, Pittsburgh, PA.

Supported in part by Grant No. DAMD 17-93-J3009 from the Department of the Army. The content of the information included in this manuscript does not reflect the position or the policy of the government and no official endorsement should be inferred.

Address reprint requests to Walter F. Good, PhD, A439 Scaife Hall, Department of Radiology, University of Pittsburgh, Pittsburgh, PA 15261-0001.

*Copyright © 1994 by W.B. Saunders Company
0897-1889/94/0703-0007\$3.00/0*

design of the quantization table required in the quantization stage of the compression process. The default quantization factors, which are provided in the standard, are not adequate for most purposes—hence, the standard encourages user-specified quantization tables. But the resulting image quality is determined by the quantization table used, so the design of these tables must be optimized for the particular compression task.

The current trend is to base quantization tables on psychophysical consideration.⁴⁻⁶ This approach is somewhat controversial when the compression is to be used for radiographic images, because psychophysical compression techniques assume prior knowledge of display characteristics and viewing conditions. The most common alternative is to use uniform quantization. The justification for this rests on the assumption that there is no a priori reason to believe that certain spatial frequencies are more important than others.

In our interpretation, the JPEG standard implies that any compression process is compatible with the standard if the encoded images it produces can be decoded with a standard JPEG decoder. Specifically, this permits various kinds of preprocessing to be applied to the input image data before the actual compression. Such preprocessing can improve the overall results of the compression/decompression cycle. For example, we have shown⁷ that when root mean square error (RMSE) is used as the measure of image fidelity, at sufficiently large compression ratios, there are significant advantages to applying a noise-removal filter to the image data before compression.

Another aspect of this preprocessing, which is somewhat unique to mammography, involves the segmentation of the image. Much of the area in a typical digitized mammogram corresponds to background pixels rather than to tissue pixels. These background pixels contain noise that is typically difficult to compress. To optimize compression in this case, it is important to minimize the amount of data required to represent these background pixels in the encoded image. This can be done in a JPEG-compatible manner by first cropping the image to the smallest rectangular area that contains all of the tissue, and then identifying the remaining background pixels and replacing them with a constant value.

With these considerations in mind, in our laboratory we have developed and begun testing software for JPEG-compatible compression of digitized mammograms.

MATERIALS AND METHODS

As described in more detail below, our compression software is consistent with the extended 12-bit version of the JPEG standard and uses a set of preprocessing routines designed specifically for mammograms. This software was written in C to be transportable, but all of the processing reported in this paper was performed on a standard 33-MHz 486 personal computer. We have tested this software by using it to compress a set of digitized mammograms at various levels of quantization, and then evaluating the quality of these compressed images by performing a just-noticeable difference (JND) study.

Preprocessing. Before applying the JPEG compression algorithm, the image data is preprocessed to improve its compressibility. This preprocessing operation is divided into two phases: first the image is segmented and cropped and the background pixel values are modified to minimize the storage requirements for the background, and then a noise-removal filter is applied to the tissue pixels.

We begin the segmentation process by automatically determining a threshold value from the histogram of the digitized source image that indicates the transition between background values and tissue values. This threshold value is used to classify pixels in the image as either tissue or background so as to produce a two-valued mask corresponding to the tissue-background separation. The thresholded image normally consists of a large contiguous region corresponding to tissue, possibly containing small holes, as well as smaller isolated clusters of points from the noisier background regions.

An initial "seed" point in the tissue area is automatically selected and grown to encompass the largest possible connected region. This produces a binary mask matching the tissue pixels in the source image. This mask is expanded by 15 pixels—ie, any pixel that is within 15 pixels of a tissue pixel is appended to the mask. This expansion is effected by convolving the binary mask with a 31×31 -pixel uniformly weighted kernel and then assigning all nonzero pixels the value 1.

We then convolve this binary mask with a 17×17 -pixel uniformly weighted kernel to smooth the edge between the mask (pixels = 1) and background (pixels = 0). This expanded and smoothed mask, whose pixel values represent an interpolation factor λ , has a value of 1 for every soft-tissue pixel and every background pixel within 7 pixels of a tissue pixel, and goes smoothly to zero over a distance of 16 pixels as the distance from tissue pixels increases. This mask and the source image are both cropped to a rectangular area that excludes as many zero pixels as possible while retaining all nonzero pixels.

A constant value, which is to be used in modifying background pixels, is determined by averaging the background pixels that are near the skin boundary. Finally, the output image is created pixel by pixel by interpolating between the source image and the constant value using the

formula

$$\text{output} = \lambda \cdot \text{source} + (1 - \lambda) \cdot \text{constant} \quad (1)$$

The resulting image has the following characteristics: (1) any 8×8 -pixel block that contains a tissue pixel will remain unchanged from the source image, eliminating the possibility that the segmentation process will alter the effect of compression on any tissue pixel; and (2) background pixels will have been modified to change in a smooth and continuous manner from their original values to the prescribed constant value as a function of distance from the nearest tissue pixel, thus reducing the possibility that the segmentation process will introduce an edge artifact at the tissue/background boundary.

The second preprocessing phase involves the application of a filter to the tissue area of the image. The particular filter tested here is a nonlinear order filter^{7,8} that was designed to modify mainly pixels with large noise components while leaving the majority of pixels unchanged. This filter operates as follows. The source image is first decomposed into a high-frequency component and a low-frequency component by unsharp masking—a process whereby a low-frequency version of the source image, obtained by smoothing the source, is subtracted from the source image to produce an edge-enhanced (high-frequency) version of the source. Pixel values in each 3×3 -pixel block of the high-frequency component are sorted. If the center pixel of the block is the highest value in the block, it is assigned the value of the next lower pixel. Similarly, if it is the lowest value, it is assigned the next higher value. Otherwise, the pixel is left unchanged. This results in a rather innocuous process of eliminating one pixel maxima or minima from the high-frequency component, while leaving most pixels unchanged. The small number of pixels changed are likely to represent noise, as their values are either greater or less than the values of all their immediate neighbors. Microcalcifications in the image are generally preserved because they rarely consist of a single pixel. Finally, the filtered high-frequency component is added back to the low-frequency component to produce the filtered source image.

These preprocessing operations require two convolutions with uniform kernels, two applications of a threshold, a region growing step, an interpolation step, and application of the nonlinear filter. In our implementation, which does not require manual intervention, all of these steps are performed in approximately one third of the time of the compression itself.

Image compression. The JPEG compression algorithm used in this study is a software package written and tested in our laboratory that implements the 12-bit version of the extended JPEG standard. For this study, Huffman encoding was used to encode the quantized coefficients and all compression ratios were based on this. The Huffman tables were derived from the statistics of each individual quantized image rather than on the ensemble statistics. Signaling information was not included in the code size when calculating compression ratios.

Results reported in this paper are based on uniform quantization of the frequency coefficients. As discussed below, any quantization scheme that preferentially degrades the high-frequency components (ie, most of the

psychophysical quantization schemes) may be contraindicated in the compression of mammographic images.

Evaluation of image-compression technique. In an effort to understand the viability of the methods described above, we carefully selected eight high-quality cases that spanned the range of image characteristics customarily encountered in mammography, with some preference being given to images containing subtle pathology. These cases were digitized and the digital data was compressed over a wide range of compression ratios by applying the above techniques.

The digitization process used a high-resolution, high-contrast sensitivity laser film digitizer (Lumisys, Sunyvale, CA) that produces a scan matrix of $4,000 \times 5,000$ pixels for an 8×10 -in film by digitizing at a $50\text{-}\mu\text{m}$ sampling interval. This pixel resolution results in a Nyquist spatial frequency of 10 cycles/mm, which preserves the signal resolution of the analog film. The modulation transfer function (MTF) at the Nyquist frequency is 30% in the fast-scan direction and 38% in the slow-scan direction. The 16-bit A/D converter of the digitizer permits a density measurement with an RMSE of less than 0.01. This represented less than one third of the noise present in the conventional film at densities exceeding 0.4.

Each digitized image was compressed at five different compression ratios. Because it is the degree of quantization rather than the compression ratio itself that determines the degradation of an image by compression, the quantization factor was used as an independent variable and not the actual compression ratio. The five quantization factors we applied were 40, 60, 80, 100, and 120, which produced average compression ratios between 15:1 and 60:1. The lowest quantization factor is at a level below which it is not expected that one could detect any effect. The highest quantization factor produced a mean compression ratio of 60:1, which is higher than what is required to permit the efficient handling of mammographic images with current digital technology, and it is in the range where studies of other algorithms applied to other image types have shown significant deterioration in image quality.^{9,10} Note that these particular quantization factors are only appropriate for images digitized at the resolution and with the noise characteristics under consideration in this study.

All images used in this study, both the compressed versions as well as the original digitized data, were printed on film with a laser film printer (Eastman-Kodak, Rochester, NY). These laser-printed films were used for all image comparisons. It was not possible for us to perform a meaningful comparison of the laser-printed images to the original (nondigitized) images because our laser printer was only capable of printing at a pitch of $\sim 80\text{ }\mu\text{m}/\text{pixel}$, and this resulted in substantial ($1.6 \times$) magnification of the laser-printed images. However, comparing two images that have both been enlarged to the same degree is possible, and if the enlargement had any effect, we believe it would have been to make it easier for readers to discriminate between noncompressed images and compressed images. The overall digitizer-processor response was calibrated to accurately duplicate the densities and contrast in the original film. Noise contributed by the laser printer is substantially less than image noise caused by quantum mottle and film granularity in the original film.

Table 1: The Effect of Image Segmentation on Tissue Block Representation (bits per block)

Image	% of Area Representing Tissue	Bits per Tissue Block (no preprocessing)	Bits per Tissue Block (after segmentation)	Improvement Factor
1	51.1	8.6	30.1	3.5
2	52.9	8.5	27.0	3.2
3	70.5	9.1	32.3	3.5
4	71.7	10.0	28.0	2.8
5	72.4	9.4	23.6	2.5
6	56.8	7.6	26.5	3.5
7	70.7	24.2	27.4	1.1
8	67.0	18.0	22.2	1.2

From the eight originals, which had each been compressed at five compression ratios, a set of 40 randomized pairs of images, each consisting of one laser-printed noncompressed image and one compressed version of the original, were assembled. Readers were asked to view the pairs side by side and were forced to choose the image they believed to have been compressed. Readers were allowed to spend as much time as desired on each pair. The reading environment was equivalent to that normally used for reading mammograms and readers were permitted to use magnifying lenses.

Altogether, 20 readers participated in these readings. They included radiologists experienced in mammography as well as physicists and engineers specializing in medical imaging.

RESULTS

The performance of our compression system is somewhat image dependent, but for the ensemble of images reported on here, we averaged about 20 s/Mbyte to preprocess and compress source data. Because decompression does not involve any steps similar to preprocessing, we were able to perform it at a rate of 15 s/Mbyte of output data. Although these speeds are not acceptable for a clinical picture archiving and communication system (PACS) environment, we believe they can be easily improved by adopting hardware optimized for the task.

A summary of the impact of image segmentation on the number of bits needed to represent tissue blocks is presented in Table 1. All images included in this table had previously been cropped to a rectangular area whose border was within 1 cm of tissue pixels in each dimension. Column 2 gives the fraction of the cropped image occupied by tissue pixels. Columns 3 and 4 give the average number of bits used to represent each 8- x 8-pixel block of tissue pixels before and after application of the segmenta-

tion process, for an overall compression ratio in both cases of 40:1. It can be seen from this table that our segmentation process increased the number of bits used to represent each tissue pixel by a factor of 2.7 on the average despite the fact that the background represented only 35.9% of the image area.

From approximately 800 observations in the JND experiment, average observers correctly detected 42%, 55%, 47%, 61%, and 75% of the images for the 15:1, 23:1, 34:1, 45:1, and 60:1 compression ratios, respectively. However, when we looked at the performance of individual readers, we found that there were two readers who were particularly good at making the discriminations required in this study. These readers were a physicist and an electrical engineer who were both specialists in image processing and were familiar with the particular artifacts caused by this type of image compression. These image-processing specialists averaged discrimination rates of 50%, 67%, 75%, 87%, and 100% for the above compression ratios. Their surprising consistency with each other and the monotonic trend of their performance led us to believe that we were actually measuring a meaningful threshold. A plot of the results of this JND study is shown in Fig 1.

DISCUSSION

Performance of compression process. The relationship between the compressibility of tissue, the compressibility of the background, and the relative areas between tissue and background

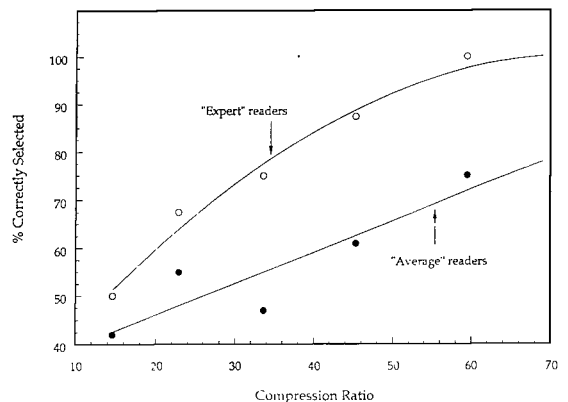


Fig 1. Performance of image-processing specialists and nonspecialists for the task of identifying compressed versus noncompressed images.

can be expressed as

$$C_t = \frac{A_t}{\frac{A_t + A_b}{C} - \frac{A_b}{C_b}}, \quad (2)$$

in which C_t and C_b are, respectively, the compression ratios for A_t and A_b , the areas corresponding to tissue and background, and C is the overall compression ratio for the image. The potential improvement factor in the number of bits per tissue pixel that can be achieved through segmentation, assuming the background is compressed to a negligible size after segmentation, is

$$\frac{1}{1 - \frac{A_b}{A_t + A_b} \cdot \frac{C}{C_b}}, \quad (3)$$

where C_b is the compression ratio of the background without segmentation. In all cases, after segmentation we were able to compress the background at a ratio of 256:1, which corresponds to using 3 bits to represent each 8- × 8-pixel block. Theoretically in some unusual cases, it would be possible to use only 2 bits per block, but this is an absolute limit for the JPEG algorithm (sequential discrete cosine transform [DCT] and Huffman encoding). The actual improvements in bits per tissue pixel are indicated in Table 1, column 5.

Note that images 7 and 8 behaved somewhat differently than the other six images. The backgrounds of these images were so dense that they were beyond the range of our digitizer. Consequently, many background pixels values were saturated at the maximum output of our digitizer, and hence, the apparent noise in the background was greatly reduced. Saturation of the background had much the same effect as our segmentation process, but has the potential to also alter tissue pixels in an undesirable way. We have included these particular images after verifying that the tissue pixels themselves were not saturated.

For images 1 through 6, our segmentation process increased the number of bits used to represent each tissue pixel by a factor of 3.2 on the average. If images 7 and 8 are also included, this effect is reduced to 2.7. The average area occupied by tissue pixels in cropped images 1

through 6 was 62.6%. It can be seen that for images in which the background is not saturated, the average improvement in bits per tissue pixel is much larger than what would be suggested simply by the relative areas occupied by the tissue pixels and background pixels.

We illustrate what happens during the segmentation process by using image number 6 from our study as a specific example. This image, which is partially shown in Fig 2, is typical of what we have observed in general. Initially, this mammogram was digitized to a 12.3-Mbyte file. Cropping reduced this by 5.7% to 11.6 Mbyte, and of this file, 56.8% of the pixels corresponded to tissue.

A plot of the compressibility of various regions of this image, which are identified in Fig 2, is presented in Fig 3. The independent variable on this graph, the quantization factor, is specified at the time of compression and directly controls the loss of information during the compression process. It can be seen that different regions of the image vary widely with respect to their compressibility and that the background is significantly less compressible than either of the tissue regions. In many cases such as this one, because the background is very noisy, it is the least compressible part of the image. When the nonsegmented version of this image was compressed at a ratio of 25:1 (0.48 bits per pixel), the compression ratio of the background was only 14.9:1 compared with 51.8:1 (0.23 bits per pixel) for the tissue pixel. The encoded file contained 0.464 Mbyte of data of which 0.337 Mbyte was needed to describe the background and the remaining 0.127 Mbyte represented tissue. In other words, despite the fact that the background was only 43% of the source image, more than 72% of the data in the encoded file was used to describe the background. After the image was preprocessed and compressed at the same 25:1 compression ratio, the constant background in the segmented image was now compressed at a ratio of 256:1 (.05 bits per pixel) and the tissue area was compressed at a ratio of 14.8:1 (0.81 bits per pixel). The background was reduced to only 0.02 Mbyte or about 4% of the encoded data. The net effect was that segmentation increased the amount of data used to represent tissue pixels by a factor of 3.5. Figure 4 compares the overall compress-




Fig 2. Image showing the three typical regions of interest in different areas representing tissue and background.

ibility of the segmented image to that of the nonsegmented image. The compressibility of a chest image is also included for comparison.

Figure 5 shows in more detail the relative number of bits per tissue pixel between the

segmented and nonsegmented images. The actual pattern of the compressibility of various parts of the image is also of interest because it provides information that eventually will be needed for optimizing the relative compression




Fig 3. Comparison of the compressibility of the three regions marked in Fig 2, with the compressibility of the image as a whole.




Fig 4. Comparison of the compressibility of mammograms with and without preprocessing to the compressibility of digitized chest images.




Fig 5. Comparison of the mean number of bits used to represent a tissue pixel for segmented versus nonsegmented images.

of different tissues. Figure 6, A and B show the pattern of bits per pixel with and without preprocessing. It is clear from a comparison of these images that preprocessing results in a dramatic shift, from bits being used to represent background to being used to represent tissue. One further observation from Fig 6B is that a large fraction of the bits used to represent tissue are in fact used to represent only the regions near skin. This observation is consistent throughout the ensemble of images we have studied to date. We believe it is mostly caused by the low compressibility of the high-frequency structures in the near-skin regions and to a lesser extent caused by the high density of the skin areas on the original films that produces increased digitizer noise. This can be seen in Fig 7, which




Fig 6. Images of bits per block before segmentation (A) and after segmentation (B) for an image compressed at 40:1.

relates the number of bits per pixel to film density for a fixed quantization level, for reasonably flat areas of an image. The increasing trend as a function of optical density is caused by the increasing noise of the digitizer as film density increases. Also shown in Fig 7 are measurements of the compressibility of skin areas as a function of density. These points are clearly above the line. We have observed that this low compressibility of the skin areas also holds for storage phosphor (computed radiography) images despite the very different noise characteristics of these images.

Future updates to the JPEG standard are likely to contain provisions for adaptive quantization. When this happens, it will be necessary to address the issue of the relative number of data bits that should be allocated to represent each particular type of tissue.

Effect of filtration. The effects of applying the nonlinear filter to mammogram number 6 is shown in Fig 8. Although the filter only changed 6.7% of the pixel values, it seemed to provide a net benefit at the higher compression ratios. Its behavior was comparable with what was achieved when a similar filter was applied to the compression of chest images.⁷ Figure 8 shows the RMSE for the compression of the filtered source along with the RMSE for the compression of the nonfiltered source as they depend on the compression ratio. It can be seen that at low compression ratios, the RMSE of the filtered image is higher because of changes in the image caused by the filter, but at higher compression




Fig 7. Plot of compressibility of different "nonskin" tissue regions as a function of film density (line). Also included for comparison are data from different points in the skin region.

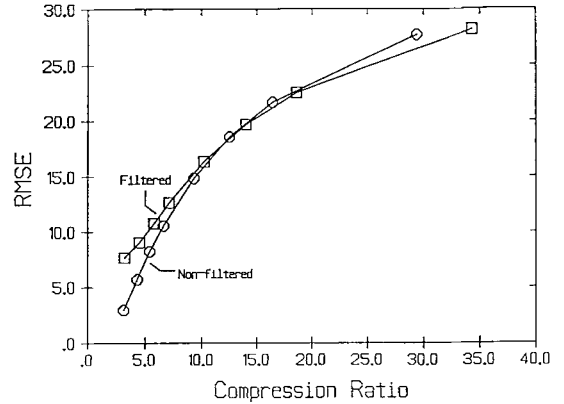


Fig 8. RMSE resulting from nonlinear filtration as a function of compression ratio. Results for both the filtered and nonfiltered image are shown.

ratios, the errors caused by compression overwhelm the error caused by filtration and become predominant. The compression-induced errors are consistently smaller for the filtered image than for the nonfiltered image, and at higher compression ratios, the filtered image actually deviates less from the source image. This behavior persisted across our image set, but was found to be image dependent. It is our view at this time that filtration of this sort is advantageous at compression ratios above 20:1, but it remains to be tested in diagnostic performance studies.

Effect of segmentation and cropping. In general, for all of our images, preprocessing the digitized data made a dramatic improvement in compressibility. This improvement can be attributed, for the most part, to the benefit of segmentation. It is actually the degree of compression of the image area corresponding to tissue pixels, or the number of bits per tissue pixel, rather than the compression ratio for the entire image array, that directly determines the impact of compression on the diagnostic quality of mammographic images. In mammography, there is wide variation from case to case in the proportion of the film occupied by tissue pixels. In all cases, after segmentation and cropping, the background will be compressed at a very high compression ratio, whereas the tissue pixels are compressed at a more modest ratio.

If images are not segmented, then cropping is particularly beneficial. However, if the images are segmented by the above procedure, the benefit of cropping is largely neutralized by the

efficiency of the segmentation process in increasing the compressibility of the background. Nevertheless, cropping can be useful, because when the image is reconstructed, it will retain its size before encoding. Maintaining excessive background pixels could have adverse effects on the time required to encode and decode the image as well as on other system components such as image displays.

Preservation of image quality. Much of a mammographic image does not contain appreciable high-frequency information, but the need to preserve microcalcifications with dimensions of less than 100 μm , which sometimes appear in mammographic films, may require that we digitize mammograms at 50 μm . This has the effect of oversampling most of the image—ie, the actual information content of the digitized mammogram is much less than the size of the image matrix would suggest.

It does not require severe quantization to achieve a high compression ratio on areas that do not contain high-frequency information. The remainder of the image, such as areas containing microcalcifications or the skin, is not as compressible, but these areas generally comprise such a small fraction of the total area that they have minimal impact on the compression ratio. Together, these factors imply that it may not be necessary or desirable to preferentially degrade the higher-frequency components in the quantization process to achieve high compression ratios, and consequently, it is possible to preserve features such as microcalcifications during the compression process.

In comparing the compressibility of mammograms to that of chest images, we found that preprocessed mammograms are very similar to chest images in compressibility; however, unprocessed mammograms are much more difficult to compress than chest images. We have included a plot in Fig 4 that was derived from a

previous study⁷ of the compression ratio of chest images versus the quantization factor. The chest image had been digitized at 100- μm resolution with 12 bits of dynamic range. At each level of quantization, the RMSE for the chest image was similar to the RMSE for mammograms compressed with the same quantization factor.

Observer performance. When the differences between images are small, it is possible that readers can detect the difference, but still not be able to tell which image is compressed. The key to differentiating between images at the lower compression ratios seemed to be the presence or absence of blocking artifact. For the compression techniques we applied in this study, blocking artifact becomes noticeable before any degradation of image resolution is obvious. Blocking artifact could be detected by image-processing specialists and nonspecialists alike, but the nonspecialist readers were more likely to misinterpret it. It is clear from discussions with some readers upon completion of their readings, that when the blocking artifact was barely visible they sometimes mistook it for actual image information and assumed the image with the artifact was the noncompressed image. The two best readers seemed to be able to avoid this mistake.

The JND study presented here is not intended to address the question of the impact of compression on diagnostic performance. Its purpose was merely to give us an opportunity to test the feasibility of the techniques we believe are appropriate for the compression of mammographic data, and to identify the range of compression ratios that merit further study. There are many possibilities for further refinements, particularly in the area of preprocessing the data, that are within the constraints of the standard, and these will likely extend the useful range of compression ratios.

REFERENCES

1. Joint Photographic Experts Group (JPEG): JPEG Technical Specification, Revision 5 (Document No. JPEG-8-R5 or JTC1/SC2/WG8 N933). ISO Central Secretariat, 1990
2. Hudson GP, Yasuda H, Sebestyen I: The international standardization of a still picture compression technique. Proc IEEE Global Telecommunications Conference & Exhibition, vol 2, November 28-December 1, 1988, pp 1016-1021
3. Wallace G, Vivian R, Poulsen H: Subjective testing

results for still picture compression algorithms for international standardization. Proc IEEE Global Telecommunications Conference & Exhibition, vol 2, November 28-December 1, 1988, pp 1022-1027

4. Daly S: Application of a noise adaptive contrast sensitivity function to image data compression. Proc SPIE 1077:217-227, 1989

5. Rabbani M, Daly S: An optimized image data compression technique utilized in the Kodak SV9600 still video transceiver. Proc. SPIE 1071:246-256, 1989

6. Ngan KN, Leong KS, Harcharan S: Cosine transform coding incorporating human visual system model. Proc SPIE 707:165-171, 1986
7. Good WF, Gur D: Quantization techniques for the compression of chest images by JPEG type algorithms. Proc SPIE 1652:114-121, 1992
8. Boncelet CG: Some uses for order statistic filtering in image compression. Proc SPIE 1247:51-57, 1990
9. Ishigaki T, Sakuma S, Ikeda M, Itoh Y, Suzuki M, Iwai S: Clinical evaluation of irreversible image compression: Analysis of chest imaging with computed radiography. Radiology 175:739-743, 1990
10. MacMahon H, Doi K, Sanada S, Montner SM, Giger ML, Metz CE, Nakamori N, Yin F-F, Xu X-W, Yonekawa H, Takeuchi H: Data compression: Effect on diagnostic accuracy in digital chest radiography. Radiology 178:175-179, 1991