Lossless Compression of Grayscale Medical Images - Effectiveness of Traditional and State of the Art Approaches

David A. Clunie
Quintiles Intelligent Imaging
521 Plymouth Road, Suite 115, Plymouth Meeting PA 19462
mailto:dclunie@idt.net

ABSTRACT

Proprietary compression schemes have a cost and risk associated with their support, end of life and interoperability. Standards reduce this cost and risk. The new JPEG-LS process (ISO/IEC 14495-1), and the lossless mode of the proposed JPEG 2000 scheme (ISO/IEC CD15444-1), new standard schemes that may be incorporated into DICOM, are evaluated here.

Three thousand, six hundred and seventy-nine (3,679) single frame grayscale images from multiple anatomical regions, modalities and vendors, were tested.

For all images combined JPEG-LS and JPEG 2000 performed equally well (3.81), almost as well as CALIC (3.91), a complex predictive scheme used only as a benchmark. Both out-performed existing JPEG (3.04 with optimum predictor choice per image, 2.79 for previous pixel prediction as most commonly used in DICOM). Text dictionary schemes performed poorly (gzip 2.38), as did image dictionary schemes without statistical modeling (PNG 2.76). Proprietary transform based schemes did not perform as well as JPEG-LS or JPEG 2000 (S+P Arithmetic 3.4, CREW 3.56). Stratified by modality, JPEG-LS compressed CT images (4.00), MR (3.59), NM (5.98), US (3.4), IO (2.66), CR (3.64), DX (2.43), and MG (2.62). CALIC always achieved the highest compression except for one modality for which JPEG-LS did better (MG digital vendor A JPEG-LS 4.02, CALIC 4.01). JPEG-LS outperformed existing JPEG for all modalities.

The use of standard schemes can achieve state of the art performance, regardless of modality. JPEG-LS is simple, easy to implement, consumes less memory, and is faster than JPEG 2000, though JPEG 2000 will offer lossy and progressive transmission. It is recommended that DICOM add transfer syntaxes for both JPEG-LS and JPEG 2000.

Keywords: Image compression, lossless compression, JPEG, JPEG-LS, JPEG 2000, DICOM.

1. INTRODUCTION

Increasingly, medical images are acquired or stored digitally. This is especially true of grayscale images that are used in radiology applications. These images may be very large in size and number, and compression offers a means to reduce the cost of storage and increase the speed of transmission.

Although the cost of storage is falling precipitously as the capacity per device increases, and the cost of transmission bandwidth is also falling; there remains a strong demand for medical image compression. Since the speed of computing is also increasing dramatically, the sophistication and complexity of compression schemes that are practical for use is increasing. For transfer over networks with high bandwidth, or for storage on electromechanical devices (disk or tape), considerable time can be spent on compression before it becomes a factor in the total transfer time.

The cost of using compression must be taken into account, however. Complex compression schemes are more costly to develop, implement, and deploy. The use of unusual or proprietary schemes has a cost (and risk) associated with the end of life of equipment (especially long term archives). It may also compromise interoperability with other equipment. The use of industry wide standards can reduce the cost and risk of the use of compression. The use of consumer industry standards (for mass produced non-medical equipment) is even better.

Much of the recent research in compression has focussed on lossy compression. Lossy compression involves deliberately discarding information that is not visually or diagnostically important. Unfortunately, lossy compression schemes may only achieve modest compression before significant information is lost. Greater compression can be achieved if some visible loss is acceptable for the clinical task. There is still controversy over the role of lossy compression for particular applications.

If only mild levels of lossy compression can be achieved for an application, then it may be that significantly improved lossless compression techniques might be more appropriate. This paper attempts to evaluate the performance of traditional and state of the art lossless compression techniques as applied to grayscale radiology images.

Emphasis is placed on those techniques that have been adopted or proposed as international standards. Particular attention is directed to the performance of the older JPEG lossless processes [1], the new JPEG-LS process [2] and the lossless mode of the proposed JPEG 2000 scheme [3]. DICOM Standard currently supports the existing JPEG processes [4]. The DICOM Committee has stated that it will not include new compression schemes in the standard unless they have applications beyond the medical imaging industry.

Lossless compression schemes can be crudely classified as follows:

- predictive schemes with statistical modeling, in which differences between pixels and their surround are computed and their context modeled prior to coding,
- transform based coding, in which images are transformed into the frequency or wavelet domain prior to modeling and coding,
- dictionary based schemes, in which strings of symbols are replaced with shorter (more probable) codes,
- ad hoc schemes (such as run length encoding).

Dictionary based schemes (such as ZIP) are widely used for text compression. Schemes for computed graphic image compression widely used on the Internet (such as GIF, TIFF LZW, and PNG) are also dictionary based. Most modern research into lossless compression involves predictive schemes with statistical modeling. The older JPEG lossless and the new JPEG-LS schemes are in this class. The lossless mode of the proposed JPEG 2000 scheme involves transformation into the wavelet domain.

This study was designed to test on a broad range of grayscale single frame medical images the hypotheses that:

1. state of the art lossless compression techniques perform substantially better than older lossless image compression techniques;
2. new international standards for compression schemes perform as well as the best state of the art lossless compression techniques;
3. state of the art lossless compression techniques perform substantially better than existing non-image based compression techniques;
4. predictive schemes with statistical modeling and transform based coding perform substantially better than dictionary based coders.

2. METHODS

The set of 3679 images tested was mostly acquired for various radiological trade shows and scientific meetings. The set was augmented by clinical images where it was lacking. The images are of multiple anatomical regions generated by multiple modalities manufactured by multiple vendors. A large collection was chosen in order to both increase the power of the study to detect small but consist trends in compression effectiveness, and to allow stratification by modality type. For most modalities, images from different vendors were pooled. Computed radiography (CR) images are distinguished from those created by other digital sensors, of which several different types and vendors were included. The mammography images were stratified by source. Optically scanned film/screen images randomly selected from the public USF database were distinguished from those acquired with digital detectors from two different vendors.

Optically scanned CT and MR laser printed film images as well as ultrasound images contain annotation burned in to the pixel data. The other images do not.

Only single frame, single component (grayscale) images were tested. Some of the image compression schemes assume that the input values are unsigned for generating difference values and for statistical modeling. Therefore any images that contained signed data were converted to unsigned (by adding the offset of the minimum value), prior to compression.

The implementations of compression schemes that were tested are listed in Table 1. All implementations (as listed in the references) were applied using their default parameters unless otherwise specified.

For 10918-1 JPEG, each of the seven alternative prediction values was tested separately. The Huffman tables were optimized for each image (two passes). Arithmetic coding was not tested since no implementation was available.

For 14495-1 JPEG-LS, the default threshold values were used. Two implementations were tested, one by this author and one by HP. In addition, a third test was performed with the author's code with the run length encoding disabled.

For CALIC and S+P, both Huffman and arithmetic entropy coders were tested.

The Packbits algorithm could only be applied to eight bit samples, but was included because it is equivalent to the RLE algorithm supported by the DICOM standard and commonly used in the Ultrasound community.

Though several alternative algorithms were initially considered for inclusion, if no robust Unix implementation was available, or memory usage was unrealistic for the larger images in the test set, they were not tested.

For the non-image based schemes that assume one symbol per byte, for those images greater than eight bits deep, both little and big endian byte orders were tested.

Compression effectiveness was evaluated by comparing the size of the compressed output with the size of the raw pixel data; header data was excluded from the calculations. Uncompressed pixel data is normally stored aligned to byte boundaries although not all bits may be used. For example, 12 bit CT data is often stored in 16 bit words with four bits of padding. The number of significant bits is usually defined in the image header, but may bear little relationship to the actual range of the data. For example, MR images usually claim to occupy the full width of 16 bit words, but may have only nine or even fewer significant bits. Accordingly, the effectiveness of compression is measured in three ways:

- relative to the uncompressed file size,
- relative to the nominal number of bits in the image,
- relative to the calculated zero order entropy of the image.

The compressed bit rate relative to the calculated zero order entropy is probably a better measure of the true effectiveness of the compression. The ratio relative to the actual file size is probably more indicative of performance in the real world.

In all cases the images were compressed, decompressed, and then compared with the original to ensure that the implementation was truly reversible. Some compression schemes failed to compress or decompress specific images, because of limitations of the scheme or because of errors in the implementation. In such cases, no compression statistics were entered for the image and scheme. The performance of other schemes on that image were included, however, in order to maximize the number of pair-wise comparisons possible.

No attempt was made to measure the speed of the implementations. All those tested were developed for research and therefore may not have been optimized for speed. Having said that, the JPEG, JPEG-LS, and SZIP codecs were noticeably faster than the others and CALIC was noticeably slower. This is to be expected given the design of the algorithms. CALIC was included in the comparison as it is the "gold standard" for the effectiveness of lossless compression, though it is considered too slow for most practical applications.

Compression performance was compared using both the parametric paired Student's *t* test, and the non-parametric Wilcoxon matched-pairs signed-ranks test. The SPSS package was used to perform the statistical analysis.

3. RESULTS

The results are summarized in Table 2, both for the whole set of image combined as well as stratified by modality type.

Table 2 compares the compressed bit-rate against the number of bits in bytes occupied by the unpacked image. Comparisons of the nominal uncompressed bit-rate and zero order entropy are not reported in the interests of brevity. The same ranking of schemes was observed regardless of which metric was used.

For all images combined, JPEG-LS and JPEG 2000 performed equally well (3.81), almost as well as CALIC (3.91). Both out-performed existing JPEG (3.04 with optimum predictor choice per image, 2.79 for previous pixel prediction as most commonly used in DICOM). Text dictionary schemes performed poorly (gzip 2.38), as did image dictionary schemes without statistical modeling (PNG 2.76). Proprietary transform based schemes did not perform as well as JPEG-LS or JPEG 2000 (S+P Arithmetic 3.4, CREW 3.56).

Stratified by modality, JPEG-LS compressed CT images (4.00), MR (3.59), NM (5.98), US (3.4), IO (2.66), CR (3.64), DX (2.43), and MG (2.62). There was considerable variation in compression depending on the type of mammogram (scanned film 2.43, digital vendor A 4.02, digital vendor B 2.43). CALIC always achieved the highest compression except for one modality. In that case, JPEG-LS performed better (MG digital vendor A JPEG-LS 4.02, CALIC 4.01). JPEG-LS outperformed existing JPEG for all modalities.

All differences observed were statistically significant at the $p<.05$ level using both parametric and non-parametric comparisons.

4. DISCUSSION

Lossless vs. lossy compression.

Despite advances in lossy compression, lossless compression remains useful for many medical imaging applications.

It is still unclear in what situations lossy compression is appropriate for short or long term archival, or for transmission for diagnostic interpretation. There is little guidance from scientific literature, professional practice standards, regulatory authorities, or the common law. The inclusion of lossy compression schemes in communications standards like DICOM does not imply that they are sanctioned for clinical use, only that the technology is available for use at the discretion of the user or the implementer.

There is no good metric by which to judge lossy schemes or determine appropriate thresholds for diagnostic use. Quantitative metrics based on analysis of the image pixels such as Mean Squared Error (MSE) and Peak Signal-to-Noise Ratio (PSNR) do not correlate well with observers’ opinions of image quality or measurement of observers’ performance. Metrics based on models of human visual perception are still in their infancy. They have not been thoroughly compared to observer performance for medical applications. Observer performance is rarely exhaustively tested since there are many potential tasks, and findings from studies of one task may not be applicable to another. The term "visually lossless" is increasingly employed for lossy compression at modest levels of loss. It may not be equivalent to "diagnostically lossless," however. Experiments based on observers’ subjective impressions that images are "diagnostically lossless" may not be a good indicator of true observer performance. Significant degrees of inter-observer and intra-observer variation on particular tasks may reduce the power of experiments on lossy compression.

For quantitative analysis, measurement of lesion size, segmentation for volume computation, computer assisted detection or characterization, images often need to be interpreted by non-human observers. Lossy compression may affect such automated or semi-automated methods. Although it may be easier to determine "quantifiably lossless" thresholds than "diagnostically lossless" thresholds, lossless compression seems likely to remain the more appropriate choice for these tasks. This may require the choice of lossless rather lossy compression for use in long-term image archives, since archived images may well be used retrospectively for unanticipated automated or quantitative applications.

"Clinical" or "region of interest" based near-lossless compression schemes exist. These identify regions of images that are determined by some criterion to be of less or no clinical interest. These regions are then discarded or selectively compressed with greater loss. Automated mechanisms to detect such regions across a broad range of image types or for specific applications (such as mammography or CT) can result in relatively high compression while maintaining truly lossless compression of the regions of interest. Some early CT compression schemes did not encode information outside the circular reconstructed area at all (perimeter encoding) and were very effective. However, if such areas are filled with a constant pixel value then most general-purpose lossless image compression schemes perform equally well.

Research into lossy image compression and encoding has produced results that are also applicable to lossless compression. Most of the recently proposed lossy schemes are transform based schemes using some form of sub-band coding. These may also be used for lossless coding if the transform is reversible, such as with an integer wavelet. The performance of the CREW, S+P, and JPEG 2000 schemes operating in lossless mode in this study illustrates the effectiveness of this approach. The rich feature set of many modern lossy schemes includes progressive transmission to fully lossless reconstruction, as well as selective encoding of regions of interest (for "clinical" near-lossless compression). The JPEG 2000 proposal includes such features [3].

Lossless compression research has also revealed alternative approaches for "near lossless" compression. The JPEG-LS scheme includes a near lossless mode in which the predictor of the pixel value may be constrained to be within a certain absolute error, instead of being exact as it is in lossless mode. This approach can achieve impressive performance without introducing the undesirable visual artifacts (such as blurring or blockiness) that are a feature of transform based lossy compression.

Types of Images That Were Evaluated

In this study, only the effect of compressing single frame grayscale images was evaluated. Many of the schemes tested may also be used to compress multi-spectral or true color images. An additional factor to evaluate is the choice of a reversible transformation into a well decorrelated color space. Compression in the RGB space is generally not very effective. Such an evaluation remains the subject of further work. Most medical applications for color compression currently make use of lossy rather than lossless techniques. The JPEG 2000 scheme includes a reversible color transformation for this application [3].

Multi-frame images, such as time based cine acquisitions of angiograms, were not studied. Many cardiac angiography applications use ISO 10918-1 JPEG lossless compression as supported in DICOM [5]. The JPEG compression is applied to each frame in isolation. No advantage is taken of interframe correlation, unlike lossy motion compression schemes such as MPEG. There is currently little interest in investigation of lossless motion compression.

Three dimensional volume images are another special case of multi-frame images. Transform based lossy and lossless compression schemes are more effective when applied in three dimensions rather than two. One study demonstrated approximately a 25% improvement for CT and MR images [6]. It is likely that Part 2 of JPEG 2000 will include support for transformation and compression in a third dimension.

Multiple frames may also be "tiled" to produce a single large image. With transform based schemes, such large composite images may compress more effectively than individual images compressed alone. This may be one reason that optically scanned sheets of laser printed CT images compress better than individual digital CT images. To take advantage of this observation, one can still use single frame image compression schemes, but take into account multiple tiles in the design of the encapsulation mechanism. A possible DICOM transfer syntax for multi-frame image objects could compose frames as tiles of larger images before compression, rather than using the current approach of compressing frames individually.

Measuring Lossless Compression Effectiveness

The effectiveness of lossless compression schemes can be described using a relative measure, "compression ratio" or by describing an absolute measure, the "bit rate" of an image. The bit rate is the average number of bits (fractional) required to encoded a pixel and is computed from the total number of bits encoded divided by the number of pixels. Such a value is useful when comparing different schemes applied to one image, or multiple images with the same bit depth, which in the case of this study they do not. Accordingly, a relative measure, the compression ratio is used here.

Compression ratios are computed from different metrics of size. One approach common in the literature is to compare the ratio of the number of bits in the uncompressed image to the number of bits in the encoded image. In the case of eight bit images common in non-medical applications this is straightforward and provides a meaningful comparison. Unfortunately, the number of bits in an uncompressed medical image may be hard to determine. Most DICOM images contain a description of bit depth that may be considered as the “nominal” bit depth, but this may be artificially large for computing compression ratios. For example, most MR images have a nominal bit depth of 16, though the actual pixel values may be encoded in fewer bits.

The true effectiveness of the compression scheme may be better indicated by comparing the encoded bit rate with a measure of how much information is "really" encoded in the image. One such measure is the "entropy" of the image, a term from information theory [7] which is essentially the average amount of information per pixel value in the image. The "zero order" entropy does not take into account any information contained in surrounding pixels. For further explanation of entropy, see [8].

A third measure of compression ratio is useful for evaluating performance in the “real world.” Images are usually stored or transmitted with pixel values aligned to byte boundaries. Eight bit image pixels are typically stored one pixel per byte. Medical images with more than 8 but fewer than 16 bits per pixel (such as 12 bit CT images) are usually stored and transmitted as two whole bytes, rather than packed more compactly. A potentially more realistic compression ratio than the others described is the comparison of the number of bytes occupied by an uncompressed image divided by the number of bytes in the compressed image. This “byte compression ratio” is used as one of the measures in this study.

Another alternative was considered but not tested. That approach is to compute the “minimum bit depth” required to encode all the pixels actually contained in an image (without recoding the pixel values). This is achieved by determining the bit width of the maximum pixel value (in the case of unsigned images). Ideally, the nominal bit depth in the image “header” would already be the minimum bit depth. The zero order entropy is a better measure of information content, however, since it is not affected by sparseness of pixel value occupation.

For the set of images tested, there was no variation in ranking of performance of compression scheme, regardless of which form of compression ratio was computed. All differences were statistically significant regardless of which metric was used. Hence, in the interests of brevity, only the byte compression ratios are tabulated in this paper.

Comparison of Compression Schemes

CALIC [9] has been reported to perform as well as, or better than, any other lossless compression scheme. It was the best performing algorithm (in terms of compression effectiveness) of all the responses to the JPEG-LS call for proposals. Its complexity, and the fact that considerably simpler schemes such as LOCO-I [10] are available, ruled it out as a contender for JPEG-LS. LOCO-I uses similar principles to CALIC, but requires only one pass through the image, and is the basis for the JPEG-LS standard. CALIC remains useful as a benchmark to which the performance of other compression schemes can be compared. As expected, CALIC with arithmetic coding performed the best, on average, for all the images in this study.

JPEG-LS and JPEG 2000 (using the integer 5,3 wavelet) tied for second best overall, with an average byte compression ratio of 3.82. Both were close in performance to CALIC with arithmetic encoding, which had an average byte compression ratio of 3.92. The use of the standard run length encoding (RLE) feature in JPEG-LS, an extension to the original LOCO-I proposal, made a noticeable difference overall, with an average byte compression ratio of 3.32 without RLE. There was no significant difference between the performances of the two different implementations of JPEG-LS that were tested, the author’s own, and that from HP.

For JPEG 2000, the integer 5,3 wavelet performed better than the other wavelet tested, the integer 2,10 wavelet, which had an average byte compression ratio of 3.67. The integer 5,3 wavelet is proposed in the committee draft as the only integer wavelet for JPEG 2000 Part 1. Other transform based schemes such as CREW [11] and S+P [12] using arithmetic coding performed well, but not as well as JPEG 2000, nor as well as JPEG-LS.

The proposed JPEG 2000 scheme out performed the other transform-based coders such as S+P and CREW, which is reassuring since many of the principles of the earlier approaches have been taken into account in the design of JPEG 2000.

The JPEG lossless scheme was evaluated using each possible choice of predictor on every image. Choosing the best predictor for each image clearly shows a dramatic improvement. The average byte compression ratio when choosing the optimum predictor for each image was 3.04 compared to 2.79 when always using selection value 1 (the previous pixel only). For a fixed choice of predictor, selection value 4 (previous pixel plus pixel above minus pixel above and left) performed best overall with an average byte compression ratio of 2.98. This improvement is not consistent across all modalities however, and in some cases other predictors performed better. These findings do call into question, however, the fixed choice of selection value 1 for some of the modality application profiles in DICOM, especially the CT/MR profile.

The performance of PNG [13] was disappointing. PNG is intended as a royalty-free replacement for the GIF format used on the Internet. The authors of PNG took the opportunity while creating a new format to add support for images greater in depth than 8 bits. The scheme makes use of the same dictionary based compressor used in zip and gzip [14]. Predictor choice is adaptively optimized on a block-by-block basis in the default mode that was used in this study. Though no attempt was made in this study to optimize the performance of PNG, it does not appear to be competitive with the other schemes tested, especially for images greater than eight bits in depth.

The SZIP scheme is a simple approach that uses the same fast entropy coder (the Rice-Golomb coder [15], [16]) as used in JPEG-LS. The absence of context modeling seems to make it less effective on the images tested. Although SZIP often performed better than JPEG lossless, it was not consistently better for all the modalities.

Previous Results

In a study of lossless compression of 3147 medical images of various modalities (CR, CT, MR, NM and US), Kivijärvi et al examined compression performance of a range of general and image specific lossless compression schemes [17]. They observed that CALIC performed consistently well (2.98). JPEG-LS did almost as well (2.81) and lossless JPEG with Huffman encoding and selection value 5 did less well (2.18). PNG also did not perform well in their test (1.90), nor did any of the general-purpose compression schemes perform as well as the image specific compression schemes. This earlier study did not have an opportunity to examine the performance of the proposed JPEG 2000 scheme. The measurement of compression ratio in their study made use of the nominal bit depth rather than byte aligned pixel data. They also measured compression and decompression time, and concluded that "CALIC gives high compression in a reasonable time, whereas JPEG-LS is nearly as effective and very fast." For example, for 512 by 512 by 16 bit images the average compression and decompression times for CALIC were 3.25 and 4.51 seconds respectively, compared with 0.85 and 0.92 seconds for JPEG-LS, and 1.91 and 1.62 seconds for lossless JPEG. Wu's implementation of CALIC, HP's implementation of JPEG-LS and the authors' own implementation of lossless JPEG was used in their study.

In another study [18], several image based and general-purpose compression schemes, including lossless JPEG, CALIC, S+P and gzip, were tested on CT, MR, PET, Ultrasound, X-Ray and Angiography images. It is not clear from the paper exactly how many images were tested, and without knowing the relative mix of modalities, it is difficult to interpret the average compression ratios specified. Only eight bit images were tested, since several of the experimental codecs used had difficulty with larger bit depth images. Larger bit depth source images were linearly scaled to maximally occupy eight bits based on the actual pixel data range. CALIC performed best with an average compression ratio of 3.65. Average compression ratios from other schemes were not explicitly specified, but extrapolating from the bar charts, S+P with arithmetic encoding achieved 3.4, lossless JPEG with Huffman encoding achieved about 2.85 and gzip achieved about 2.05. The paper also tested several other general purpose (STAT) and image based (Binary Tree Predictive Coding - BTPC) schemes that performed well, but not as well as CALIC.

These results are entirely consistent with what has been observed in the present study. The earlier papers tested some schemes not included in this study. None of them appear to have performed well enough to warrant use as a replacement for CALIC as a benchmark, or to be considered as competitors to the new standard schemes (JPEG-LS and JPEG 2000).

This study also examines the performance of lossless compression schemes on images acquired with new detector types not available to previous authors. These include flat-panel and other novel digital X-ray sensors used for both general-purpose projection radiography and mammography.

JPEG-LS has also been tested for use in combination with other approaches to improve compression effectiveness. One study used compression of the dynamic range of CR images, followed by lossless or near-lossless encoding using JPEG-LS, then re-expansion of the dynamic range [19]. This approach achieved compression ratios on CR images of 6:1 without visually perceived loss. This study also commented on the effectiveness of the run length mode in JPEG-LS for compressing background outside the exposed field.

The present study does not address the issue of lossless compression of coronary angiographic images. This is an important class of images to consider, since DICOM CD-R has become the medium of choice to replace 35 mm cine film in this application. A recent paper describes the performance of an experimental integer wavelet based scheme similar to what is proposed in JPEG 2000 on eight bit coronary angiograms [20]. One of the algorithms described achieved an average compression ratio of 4.20, about the same as CALIC. Lossless JPEG achieved only 3.06 on the same images, though it is not clear which predictor was selected.

5. CONCLUSIONS

The results of the experiments described in this paper confirmed the validity of the hypotheses that:

1. state of the art lossless compression techniques perform substantially better than older lossless image compression techniques;
2. new international standards for compression schemes perform as well as the best state of the art lossless compression techniques;
3. state of the art lossless compression techniques perform substantially better than existing non-image based compression techniques;
4. predictive schemes with statistical modeling and transform based coding perform substantially better than dictionary based coders

The use of standard schemes can achieve state of the art performance, regardless of modality. JPEG-LS is simple, easy to implement, consumes less memory, and is faster than JPEG 2000, though JPEG 2000 supports progressive transmission.

It is recommended that DICOM consider the adoption of transfer syntaxes for both JPEG-LS, as well as JPEG 2000. Both offer considerably improved lossless compression performance over the JPEG transfer syntaxes currently in the standard. The adoption of new standard methods will reduce the proliferation of private transfer syntaxes that compromise interoperability. It will also reduce the use of non-DICOM protocols that support better compression schemes.

6. ACKNOWLEDGEMENTS

The author thanks those individuals, institutions and vendors who have contributed to the collection of images used. In particular, Brad Erickson (Mayo Clinic), Andrew Maidement and Michael Albert (Thomas Jefferson University), John French, Charles Parisot, Sandrine Bruneau, Serge Muller and Herve Hoehn (GE Medical Systems), and Mark Huffe (Sterling) all contributed images specifically for this study. Gadiel Seroussi of HP Laboratories provided answers to questions relating to the implementation of JPEG-LS, and Martin Boliek of Ricoh provided the CREW test code. This study would not have been possible if compression researchers did not routinely place their code and papers on the Internet for public access.

7. REFERENCES

1. ISO/IEC 10918-1 *Digital image compression and coding of continuous-tone still images.*
2. ISO/IEC 14495-1 *Lossless and near-lossless and coding of continuous-tone still images (JPEG-LS).*
3. ISO/IEC JTC1/SC29/WG1. *WG1N1523 JPEG 2000 Part I Committee Draft Version 1.0.* 1999.
4. DICOM (Digital Imaging and Communications in Medicine), *PS 3.5 Data Structures and Encoding*, NEMA, Rossyln VA, 1999. "http://medical.nema.org/dicom/1999/draft/99_05dr.pdf"
5. DICOM (Digital Imaging and Communications in Medicine), *PS 3.11 Media Application Profiles*, NEMA, Rossyln VA, 1999. "http://medical.nema.org/dicom/1999/draft/99_11dr.pdf"
6. Bilgin A, Zweig G, Marcellin MW, "Efficient lossless coding of medical image volumes using reversible integer wavelet transforms", *Proc. 1998 Data Compression Conference*, Storer J, Cohn M eds, IEEE Computer Society Press, Los Alamitos CA, 1998. "http://www-spacl.ece.arizona.edu/Publications/Papers/C98_2.ps"
7. Shannon CE, Weaver W, *The mathematical theory of communication*. University of Illinois Press, Urbana IL, 1949.
8. Rabbani M, Jones PW, *Digital image compression techniques*, SPIE Press, Bellingham WA, 1991.
9. Wu X, Memon N, "Context-based, adaptive, lossless image codec", *IEEE Trans. On Communications*, **45**, pp. 437-444, 1997.
10. Weinberger MJ, Seroussi G, Sapiro G, "LOCO-I; A low complexity, context-based, lossless image compression algorithm", *Proc. 1996 IEEE Data Compression Conference*, Storer J, Cohn M eds., pp 140-149, IEEE Computer Society Press, Los Alamitos CA, 1996. "http://www.hpl.hp.com/loco/dcc96copy.pdf"
11. Zandi A, Allen JD, Schwartz EL, Boliek M, "CREW: Compression with reversible embedded wavelets", *Proc 1995 IEEE Data Compression Conference*, , Storer J, Cohn M eds., IEEE Computer Society Press, Los Alamitos CA, 1995.

12. Said A, Pearlman W, "Reversible image compression via multi-resolution representation and predictive coding", *SPIE VCIP Symposium*, Cambridge, MA, 1993. "ftp://ipl.rpi.edu/pub/EW_Code/SpieVcip93.ps.gz"
13. Randers-Pehrson G, et al, *PNG (Portable Network Graphics) specification version 1.1*. PNG Development Group. February 1999. "http://www.cdrom.com/pub/png/spec/". (Also ISO/IEC JTC1/SC24 ISO/IEC DIS 15948).
14. Deutsch P, Gailly JL, *ZLIB Compressed data format specification version 3.3*. Internet RFC 1950. Aladdin Enterprises. May 1996. "ftp://ds.internic.net/rfc/rfc1950.txt".
15. Rice RF, *Some practical universal noiseless coding techniques*, Technical Report JPL-79-22.Pasadena CA 1979.
16. Golomb SW, "Run-length encodings", *IEEE Transactions on Information Theory*, 1966 pp 140-149.
17. Kivijärvi J, et al, "A comparison of lossless compression methods for medical images", *Computerized Medical Imaging and Graphics*, **22**, pp 323-339, 1998.
18. Denecker K, Van Overloop J, Lemahieu I, "An experimental comparison of several lossless image coders for medical images", *Proc. 1997 IEEE Data Compression Conference*, Storer J, Cohn M eds., IEEE Computer Society Press, Los Alamitos CA, 1997. "http://www.elis.rug.ac.be/~denecker/doc/papers/dcc97/workshop.ps.gz"
19. Young SS, Whiting BR, Foos DH, "Statistically lossless image compression for CR and DR", *Proceedings of SPIE: Medical Imaging 1999: Image Display*, **3658**, SPIE Press, Bellingham WA, 1999.
20. Munteanu A, Cornelis J, Cristea P, "Wavelet-based lossless compression of coronary angiographic images", *IEEE Transactions on Medical Imaging*, **18**, pp 272-281, 1999.

Table 1 – Compression Schemes and Implementations Tested.

Algorithm	Type	Implementation	Version	Source
JPEG Lossless (ISO 10918-1)	Predictive	Stanford PVRG	1.2 1993/02/24	ftp://havefun.stanford.edu/pub/jpeg
JPEG LS (ISO 14495-1)	predictive+rle	HP LOCO-I	0.90	http://www.hpl.hp.com/loco/software.htm
JPEG LS (ISO 14495-1)	predictive+rle	Author's	1999/02/11	http://idt.net/~dclunie/jpegls.html
JPEG LS (ISO 14495-1) – no runs	predictive	Author's	1999/02/11	http://idt.net/~dclunie/jpegls.html
JPEG 2000 with integer 5,3 wavelet	transform	WG1SC29 VM	4.2	-
JPEG 2000 with integer 2,10 wavelet	transform	WG1SC29 VM	4.2	-
SZIP NASA's Rice coder	predictive+rle	UMN USES	1.4 1996/08/1	ftp://groucho.mrc.unm.edu/pub/szip/
CALIC Huffman	predictive	Wu and Memon	1995	ftp://ftp.csd.uwo.ca/pub/from_wu/v.huff
CALIC Arithmetic	predictive	Wu and Memon	1995	ftp://ftp.csd.uwo.ca/pub/from_wu/v.arith
S+P Huffman	predictive	Said and Pearlman	4.04 1995/06/15	ftp://ipl.rpi.edu/pub/EW_Code/
S+P Arithmetic	predictive	Said and Pearlman	4.04 1995/06/15	ftp://ipl.rpi.edu/pub/EW_Code/
TIFF Packbits	rle	Loeffler's libtiff	3.4	ftp://ftp.uu.net/graphics/tiff/
CREW	transform	Ricoh	0.4.1 1996/10/20	-
PNG	dictionary	pnmtopng +libpng +libgzip	xxx +0.89c +1.03	ftp://swrinde.nde.swri.edu/pub/png/applications/ ftp://swrinde.nde.swri.edu/pub/png-group/src/history/libpng-0.89c.tar.gz ftp://ftp.uu.net/pub/archiving/zip/zlib/
LZ77 (8 bit symbols)	dictionary	Unix gzip	1.2.4	ftp://gatekeeper.dec.com/pub/GNU/gzip-1.2.4.tar.Z
LZW (8 bit symbols)	dictionary	Unix compress	Solaris 5.x	-
Byte Adaptive Huffman (8 bit)	dictionary	Unix pack	Solaris 5.x	-

Table 2 – Byte Compression Ratio compared with Compression Scheme and Modality – Full Image Set.

	Compression Ratio (Raw Image Bytes vs. Compressed Image Bytes) (top five (or more if ties) for each modality are in bold italics)																			
Compression Scheme (1)	All	<= 8	> 8	CT	CT (film)	MR	MR (film)	NM	US	All RG	DX	CR	All MG	MG (film)	MG DX1	MG DX2	RF	XA	PX	IO
Number of Images	3679	731	2948	1180	44	1331	12	71	268	143	49	76	81	51	10	20	127	35	18	212
PACKBITS (2)	1.66	1.66	0	0	0	0	0	2.58	1.79	1.21	1.21	0	0	0	0	0	1.33	0	1.32	1.19
Unix pack	1.73	1.63	1.75	1.6	1.32	1.88	1.41	3.13	1.86	1.36	1.29	1.43	1.32	1.3	1.63	1.21	1.35	2	1.1	1.3
Unix compress BE	2.24	2.47	2.18	2.08	1.65	2.14	1.65	4.77	2.47	1.92	1.68	2.09	1.42	1.33	2.13	1.28	1.94	***3.73***	2.17	1.93
Unix compress LE	2.24	2.47	2.18	2.08	1.65	2.14	1.65	4.77	2.47	1.92	1.68	2.09	1.42	1.33	2.13	1.28	1.94	***3.73***	2.17	1.93
GNU gzip BE	2.38	2.54	2.35	2.24	1.72	2.3	1.72	5.36	2.53	1.99	1.77	2.18	1.67	1.52	2.65	1.54	1.96	3.48	2.34	1.96
GNU gzip LE	2.39	2.54	2.35	2.25	1.72	2.3	1.72	5.36	2.53	1.99	1.77	2.18	1.67	1.52	2.65	1.54	1.96	3.49	2.34	1.96
JPEG SV 3	2.58	2.14	2.69	2.67	2.07	2.62	2.14	4.44	1.98	2.53	2.13	2.87	2.18	2.13	3.02	1.91	2.16	2.86	2.2	1.93
JPEG SV 2	2.76	2.34	2.87	2.86	2.2	2.8	2.27	4.68	2.13	2.65	2.17	3.02	2.26	2.15	3.21	2.07	2.35	3.03	2.56	2.11
PNG	2.76	3.31	2.62	2.64	1.7	2.54	1.77	5.07	3.04	2.18	1.84	2.28	1.89	1.69	3.14	1.76	2.65	2.6	3.78	2.44
JPEG SV 1	2.79	2.37	2.89	2.91	2.19	2.8	2.29	4.76	2.3	2.64	2.2	2.99	2.31	2.26	3.2	1.96	2.24	3.11	2.56	2.1
JPEG SV 7	2.85	2.4	2.96	2.99	2.24	2.89	2.35	4.42	2.22	2.73	2.24	3.1	2.32	2.26	3.14	2.05	2.39	2.96	2.7	2.14
S+P Huffman	2.87	1.95	3.11	3.19	1.73	3.15	1.86	4.47	1.52	2.16	1.79	2.31	1.85	1.85	0	0	2.47	2.96	1.71	1.51
JPEG SV 6	2.89	2.42	3	3.09	2.3	2.91	2.41	4.2	2.26	2.64	2.14	3.02	2.28	2.19	3.16	2.1	2.31	2.91	2.92	2.12
JPEG SV 5	2.92	2.45	3.03	3.13	2.26	2.92	2.38	4.5	2.33	2.64	2.16	3.01	2.3	2.24	3.16	2.02	2.27	2.94	2.89	2.13
JPEG SV 4	2.98	2.53	3.09	3.27	2.3	2.91	2.44	4.44	2.36	2.56	2.07	2.94	2.25	2.15	3.18	2.05	2.19	2.99	3.32	2.16
JPEG best	3.04	2.62	3.14	3.29	2.31	2.97	2.45	4.86	2.38	2.76	2.26	3.14	2.38	2.32	3.21	2.11	2.4	3.14	3.41	2.3
NASA szip	3.09	2.81	3.17	3.17	2.24	3.12	2.41	5.36	2.72	2.79	2.24	3.21	2.42	2.26	3.68	2.19	2.5	3.15	2.81	2.15
JPEG-LS no run	3.31	2.98	3.39	3.61	***2.56***	3.22	***2.77***	4.56	2.86	2.99	2.39	3.41	2.52	2.39	3.43	2.37	2.8	3.26	3.7	2.52
S+P Arithmetic	3.4	2.46	3.64	3.72	1.79	***3.71***	1.97	5.78	1.68	2.28	1.85	2.43	1.93	1.93	0	0	2.86	3.49	1.88	1.62
CREW	3.56	3.38	3.6	3.67	***2.54***	3.56	2.76	5.32	2.91	***3.15***	***2.45***	***3.62***	***2.61***	***2.41***	***3.82***	2.4	2.9	3.42	3.99	***2.53***
JPEG2000 2x10	3.66	3.43	***3.72***	***3.83***	2.5	***3.67***	2.73	5.38	2.93	3.12	2.41	3.61	2.54	2.38	3.78	2.4	2.91	3.39	***4.02***	2.51
CALIC Huffman	***3.67***	***3.52***	3.71	3.79	***2.54***	***3.59***	2.74	***6.05***	***3.2***	3.12	2.41	3.61	***2.6***	***2.4***	***3.94***	***2.42***	***3***	***3.51***	3.52	***2.57***
JPEG-LS HP	***3.81***	***3.86***	***3.8***	***4***	***2.58***	***3.59***	***2.77***	***6***	***3.39***	***3.15***	***2.43***	***3.64***	***2.62***	***2.43***	***4.02***	***2.43***	***3.08***	***3.55***	***4.2***	***2.66***
JPEG-LS MINE	***3.81***	***3.86***	***3.8***	***4***	***2.58***	***3.59***	***2.77***	***5.98***	***3.4***	***3.15***	***2.43***	***3.64***	***2.62***	***2.43***	***4.02***	***2.43***	***3.08***	***3.55***	***4.2***	***2.66***
JPEG2000 5x3	***3.81***	***4.11***	***3.74***	***3.88***	2.49	***3.63***	***2.72***	***5.69***	***2.97***	***3.14***	***2.43***	***3.63***	2.4	2.38	3.53	***2.41***	***2.94***	3.42	***3.95***	2.51
CALIC Arithmetic	***3.91***	***3.93***	***3.9***	***4.01***	***2.59***	***3.72***	***2.8***	***7.04***	***3.5***	***3.23***	***2.46***	***3.74***	***2.64***	***2.45***	***4.01***	***2.45***	***3.14***	***3.62***	***4.22***	***2.7***

Notes: 1. Sorted in order of increasing compression ratio for all images combined.
2. The PACKBITS implementation only supports 8 bit pixels so not all modalities could be tested and data present reflects only the subset of images for each modality that were <= 8 bits in depth