LOSSLESS COMPRESSION OF COLOR MEDICAL RETINAL IMAGES

Roman Starosolski
Institute of Computer Science
Silesian University of Technology
Akademicka 16, 44-100 Gliwice, Poland
e-mail: Roman.Starosolski@polsl.pl

Gerald Schaefer
School of Computing and Informatics
Nottingham Trent University
Nottingham, NG11 8NS, United Kingdom
e-mail: Gerald.Schaefer@ntu.ac.uk

KEYWORDS
Lossless image compression, retinal images, color space transforms, image compression standards.

ABSTRACT

We analyzed the performance of several lossless image compression algorithms, namely Lossless JPEG, JPEG-LS, JPEG2000, PNG, and CALIC on a large set of medical retinal images organized into groups according to capturing conditions and retinal region. Overall we found JPEG-LS to be the best performing algorithm when we consider both compression speed and compression ratio. Although we observed some variance among groups JPEG-LS consistently obtains both the speed and the ratio close to the best of all tested algorithms. We also found that the compression ratios may be further improved by applying (reversible) color space transforms and histogram packing transforms to the retinal images before compressing them. However, for certain images color transforms may produce worse compression ratios.

INTRODUCTION

Diabetic retinopathy is the leading cause of blindness in the adult population. In order to effectively identify patients suffering from the disease, mass-screening efforts are underway during which digital images of the retina are captured and then assessed by an ophthalmologist. Since, in order to identify features such as exudates and microaneurysms, which are typically small in extent, retinal images have to be captured at high resolutions which in turn produce images of large file size. Therefore, in order to reduce the demands on resources such as storage space and bandwidth, compression of the images seems necessary.

In this paper we set out to identify a suitable compression algorithm for retinal images. Such an algorithm, in order to be useful in a real-life medical picture archiving and communications system (PACS), should not only reduce the size of images significantly but has to be fast enough as well. Also, it should be covered by a widely recognized industry standard. We analyze the performance of standard or quasi-standard lossless image compression algorithms, namely Lossless JPEG, JPEG-LS, JPEG2000, PNG, and CALIC, on a large set of medical retinal images.

Since retinal images are (RGB) color images and since for certain images individual color channels have relatively sparse histograms, we experiment with some reversible transforms, namely color space conversion and histogram packing transforms that may be used to improve compression ratios.

REVERSIBLE TRANSFORMS

Color Space Conversions

It is known that the Red, Green and Blue channels of the RGB color space are highly correlated. To remove this inter-channel correlation one could use the optimal KLT/PCA transform, but as the KLT coefficients are calculated for each image individually based on the analysis of the whole image the use of KLT transform would slow down the compression process noticeably due to the high computational complexity involved. Therefore, in lossy compression algorithms like JPEG or JPEG2000, a standard transform to a different color space such as YCbCr transform is employed. In the case of lossless coding, clearly a reversible transform must be employed to retain all image information. In this paper we report results of applying two such transforms.

Reversible Multiple Component Transformation (RMCT) of the JPEG2000 standard core coding system (ISO/IEC and ITU-T 2002b) is an integer approximation of the YCbCr transform defined as

$$\begin{cases} C_1 = R - G \\ C_2 = B - G \\ Y = \lfloor (R + B + 2G)/4 \rfloor \end{cases} \tag{1}$$

where R, G, and B are the pixel values of the red, green, and blue channels respectively. The disadvantage of this transform is that it expands the dynamic range of the transformed coefficients, i.e. if R, G, and B are of 8-bit precision (as often the case) then C_1 and C_2 require 9-bit representation and Y needs 8 bits. We report the effects of applying this transform on the results obtained by the JPEG2000 compression algorithm.

Reversible Color Transform (RCT) of the JPEG-LS extended standard (ISO/IEC and ITU-T 2002a) is a modulo-arithmetic version of the RMCT defined as

Proceedings 20th European Conference on Modelling and Simulation
Wolfgang Borutzky, Alessandra Orsoni, Richard Zobel © ECMS, 2006
ISBN 0-9553018-0-7 / ISBN 0-9553018-1-5 (CD)

$$\begin{cases} C_1 = (R - G + 2^{N-1}) \bmod 2^N \\ C_2 = (B - G + 2^{N-1}) \bmod 2^N \\ Y = (G + \lfloor (C_1 + C_2)/4 \rfloor - 2^{N-2}) \bmod 2^N \end{cases} \quad (2)$$

where N is the color channel bit depth. Due to modulo clipping this transform may introduce rapid changes of intensities in channels of transformed color space even if the intensities in RGB space were changing gradually. However, it does not expand the dynamic range. We report effects of applying the RCT transform on results obtained by all of the tested algorithms.

Histogram Packing

Another reversible transform that may improve compression ratios of certain images is histogram packing. In some images the actual number of active pixels' intensity levels is smaller than implied by the nominal bit depth. Furthermore, active levels do not occupy the continuous subinterval of the nominal range; instead they are distributed more or less sparsely throughout a part or all the nominal intensity range. We call such images as having sparse histogram. Image compression algorithms are based on sophisticated assumptions as to characteristics of the images they process. Sparse intensity level histograms are different from what is expected by lossless image compression approaches, both in case of predictive and of transform coding.

To improve the compression ratios of such images we apply a histogram packing transform (Pinho 2001; Ferreira and Pinho 2002). An off-line histogram packing simply maps all the active levels to the lowest part of the nominal intensity range (order-preserving one-to-one mapping). In case of color images the transform is applied to each channel individually. Reversing the off-line packing requires the information describing how to expand the histogram after decompressing an image and hence this data has to be encoded along with the compressed image. There are various methods of encoding this information (Starosolski 2005), we employed the Bit-Array method where for all nominally available levels we encoded the information whether the specific level is active. The overhead associated with histogram packing (which was taken into account in calculating bit rates) is hence 96 bytes (= 3 x 256 bits) for full color images. For obvious reasons, if both histogram packing and color space conversion are employed, the histogram transformation is applied prior to the color space transform.

RETINAL IMAGE DATASET

A large set of medical retinal images containing images obtained at various centers was used in our experiments. The set is large both with regards to the number of images as well with regards to the actual sizes of individual images. Images contain between 1.4 and 3.5 millions pixels and hence require, since they are 24-bit RGB color images, between 4 and 10 MB of storage space.

In order to check whether there are certain images that are especially susceptible to compression (or especially hard to compress) our set was divided into groups in two different ways (as detailed in Table 1) according to the following criteria:

- Retinal region: the whole set is divided into three groups (*Nasal*, *Posterior*, and *Temporal*) according to this criterion.
- Capturing conditions: a criterion we introduced based on certain non-diagnostic properties of images (like the image size). According to this criterion, we divided the whole set into five groups (*A*, *B*, *C*, *D*, and *E*). In each group, all images are of the same size. Since group *C* contains only 3 images we decided to not report results for this group. We note however, that each of the groups *Nasal*, *Posterior*, and *Temporal* contains one image from group *C*, and when reporting results for all the images (*All*) we include results of these 3 images.

Table 1: Test dataset of retinal images

Group	No of images	Average Size [pixels]
Nasal	252	2214864
Posterior	301	2203523
Temporal	250	2213720
A	162	1382405
B	303	1382405
C	3	1504006
D	282	3333126
E	53	3538950
All	803	2210257

EXPERIMENTAL PROCEDURE

Experimental results were obtained on a HP Proliant ML350G3 computer equipped with two Intel Xeon 3.06 GHz (512 KB cache memory) processors and Windows 2003 operating system. Single-threaded applications of algorithms used for comparisons were compiled using Intel C++ 8.1 compiler. To minimize effects of the system load and the input-output subsystem performance the implementations were run several times; the time of the first run was ignored while the collective time of other runs (executed for at least one second, and at least 3 times) was measured and then averaged. The time measured is the sum of time spent by the processor in application code and in kernel functions called by the application, as reported by the operating system after application execution. The compression speed is reported in megabytes per second [MB/s], where $1\text{MB} = 2^{20}$ bytes. We note that we actually measure the speed of the specific implementation of the given algorithm on the particular computer system, not the absolute speed of the algorithm itself. In order to make the results obtained useful in the real life, we utilized popular and "standard" (if available) implementations. The computer system we used, in terms of the processor speed, amount of the installed memory etc., is similar to currently manufactured PCs. The compression ratio we give in bits per pixel [bpp] $8\,e\,/\,n$, where e is the size in bytes of the compressed image including the header and n is the number of pixels in the image. We note that smaller ratios mean better compression and that uncompressed images are stored using 24 bpp.

COMPRESSION ALGORITHMS

In this section, we characterize, very briefly, the algorithms and implementations we used in this study:

- Lossless JPEG: the former JPEG committee standard for lossless image compression (Langdon et al. 1992). The standard describes predictive image compression algorithm with Huffman (Huffman 1952) or arithmetic (Moffat *et al.* 1998) entropy coding. We used the Cornell University implementation by Huang and Smith, (version 1.0, ftp://ftp.cs.cornell.edu/pub/multimed/ljpg.tar.Z) which uses Huffman codes. The results are reported for the predictor function SV 7 which gave the best average compression ratio on the test dataset.
- JPEG-LS: a standard of the JPEG committee for lossless and near-lossless compression of still images (ISO/IEC and ITU-T 1999). The standard describes low-complexity predictive image compression algorithm with entropy coding using modified Golomb-Rice (Golomb 1966; Rice 1979) family. The algorithm is based on the LOCO-I algorithm (Weinberger *et al.* 1996; Weinberger *et al.* 2000). We used the Signal Processing and Multimedia Group, University of British Columbia implementation (version 2.2, ftp://ftp.netbsd.org/pub/NetBSD/packages/distfiles/jpeg_ls_v2.2.tar.gz).
- JPEG2000: a fairly recent JPEG committee standard describing algorithm based on wavelet transform image decomposition and arithmetic coding (ISO/IEC and ITU-T 2002b). Apart from lossy and lossless compressing and decompressing of whole images it delivers many interesting features (progressive transmission, region of interest coding, etc.) (Christopoulos et al. 2000; Adams 2001). We used the JasPer implementation by Adams (version 1.700.0, http://www.ece.uvic.ca/ ~mdadams/jasper/).
- PNG: a standard of the WWW Consortium for lossless image compression (WWW Consortium 1996). PNG is a predictive image compression algorithm, using the LZ77 (Ziv and Lempel 1977) algorithm and the Huffman codes. We used the pnmtopng implementation by Lehmann, van Schaik, and Roelofs, (version 2.37.6, a part of the NetPBM 10.25 toolkit, http://netpbm.sourceforge.net) compiled with libraries LibPNG 1.2.8 (http:// libpng.sourceforge.net) and ZLIB 1.2.2 (http:// www.gzip.org/zlib/). In PNG we have a choice of the predictor functions (filters) used to decorrelate the images. The results are reported for the filter selected by the encoder, from among all 4 available ones, for each image individually. In most cases either the "avg" or the "paeth" filter was best for a given image; employing automatic selection of filter decreases the compression speed by about 10% compared to compression performed using fixed "avg" or "paeth" filter.
- CALIC: a relatively complex predictive image compression algorithm using arithmetic entropy coder, which because of its usually good compression performance is commonly used as a reference for other image compression algorithms (Wu and Memon 1997; Wu 1997). We used implementation by Yuan. In contrast to the other algorithms CALIC is designed for grayscale images only. The CALIC speed is reported based on total time of compressing individual color channels of a given image (ratio based on total size of compressed color channels).

EXPERIMENTAL RESULTS AND DISCUSSION

In Tables 2 and 3 we report the compression results obtained on the image dataset without applying the reversible transforms described earlier. Results are given in terms of compression speed (Table 2) and compression ratio (Table 3), for each of the groups in the dataset as well as average results for the whole set. The numbers are calculated as an average for all images contained in the group; since not all groups contain the same number of images the average results for all images may be slightly different from the average of all groups.

Table 2: Compression speed results [MB/s]

Group	**PNG**	**Lossless JPEG**	**JPEG-LS**	**JPEG2000**	**CALIC**
Nasal	2.9	16.4	14.7	3.3	2.7
Posterior	2.9	16.5	14.7	3.2	2.7
Temporal	2.9	16.4	14.6	3.3	2.7
A	3.4	17.3	15.9	3.7	2.9
B	2.5	16.1	15.0	3.3	2.7
D	3.1	16.4	13.6	3.3	2.6
E	2.8	16.7	14.7	1.5	2.7
All	2.9	16.5	14.7	3.3	2.7

For any specific algorithm it can be seen that there are no significant differences between groups defined according to retinal region (*Nasal*, *Posterior*, and *Temporal*), also the variability indexes of results in these groups (e.g. from 13% to 16% in case of JPEG-LS ratio) are close to variability indexes for the whole set (15% for JPEG-LS). Noticeable differences between groups, especially in terms of compression ratio, and lower variabilities within specific group are observed in the case of groups defined according to image size and capturing conditions (11% for group *D*, 2%–3% for *A*, *B*, and *E*). Since similar characteristics are observed for results for images after applying reversible transforms described earlier, in the reminder of this paper we report average results for the whole set and for the *A*, *B*, *D*, and *E* groups only.

Table 3: Compression ratio results [bpp]

Group	PNG	Lossless JPEG	JPEG-LS	JPEG2000	CALIC
Nasal	7.580	8.123	6.830	7.144	6.610
Posterior	7.530	8.087	6.791	7.113	6.559
Temporal	7.539	8.098	6.840	7.160	6.629
A	6.157	6.632	5.255	5.784	5.154
B	8.539	8.602	7.151	7.403	6.829
D	7.330	8.499	7.446	7.703	7.243
E	7.249	7.515	6.239	6.636	6.124
All	7.549	8.102	6.819	7.138	6.597

Considering the compression speed, Lossless JPEG and JPEG-LS obtain similar high speeds of about 15 MB/s. The speeds of the remaining algorithms are roughly 5 times lower. While our aim in this study was to analyze standard and quasi-standard image compression algorithms we note that there exist other methods with compression speeds of over 2 times that of Lossless JPEG and JPEG-LS (Consultative Committee for Space Data Systems 1997; Starosolski 2006). Among the slower algorithms, CALIC and JPEG2000 use an arithmetic entropy coder in contrast to PNG which is based on faster techniques (LZ77 and the Huffman coding). Based on this, one could expect PNG to obtain compression speeds close to Lossless JPEG with Huffman coding. Therefore, the low speed of PNG might be due to the particular implementation used (similar results were obtained a set of medical thermal images (Schaefer *et al.* 2005)). For JPEG2000 we notice that the speed for images of size greater than about 10MB (group *E*) is much lower compared to speed for smaller images (including images of sizes close to 10MB); possibly the implementation used inefficiently manages images of sizes exceeding a certain threshold.

Inspecting the results in terms of compression ratios we can see that the compression performance of Lossless JPEG is consistently the worst among all tested algorithms while the ratio of JPEG-LS is the second best, both overall (CALIC is better by about 3%) and for nearly all the groups as well (in case of the *D* group PNG is between CALIC and JPEG-LS).

JPEG-LS hence seems to be the best choice for practical compression of untransformed color retinal images; using this algorithm we can efficiently reduce the storage requirements by a factor of about 3.5. Employing lossless image compression such as JPEG-LS in a medical picture archiving and communications system (PACS) will hence significantly increase the mass storage capacity and at the same time decrease transmission times within the existing infrastructure which in turn can speed up clinician's work. In addition, JPEG-LS is covered by recent international standards including in particular the DICOM standard (National Electrical Manufacturers Association 2001) for storing medical imagery which is widely recognized by vendors of medical devices and medical software.

Next, we examined the application of reversible transforms as described above. Subjecting images to color space conversion or histogram packing prior to compression improves the compression speed of transformed images only insignificantly (by a few %), hence for brevity we do not report detailed speed results here. Compression ratios for images compressed after applying color transforms are reported in Table 4. Except for the column marked "JPEG2000*" which gives the JPEG2000 ratios after applying the JPEG2000 RMCT transform from Equation 1, we report ratios for images after applying the RCT transform from Equation 2.

Table 4: Compression ratios after color transform [bpp]

Group	PNG	Lossless JPEG	JPEG-LS	JPEG2000*	JPEG2000	CALIC
A	5.713	6.137	4.567	5.177	5.329	4.304
B	8.997	9.179	7.804	7.400	7.795	7.595
D	6.971	7.631	6.341	6.706	6.841	6.105
E	6.785	6.895	5.473	5.981	6.071	5.282
All	7.474	7.869	6.481	6.613	6.849	6.253

Applying color transforms does not change the relative performance among algorithms tested. Interestingly, the RCT does not improve compression ratios in all cases, applying it to group *B* images actually worsens the ratios of all algorithms by 5% to 10%. For the other groups it improves ratios by 5% to 15%. Greater differences are observed for more sophisticated algorithms (JPEG-LS, JPEG2000, and CALIC). On average, applying RCT improves ratios; for JPEG-LS, JPEG2000, and CALIC the improvement is about 5%. Comparing JPEG2000 RCT and RMCT ratios it appears that the latter transform is more effective for lossless compression of retinal images. It does not worsen (negligibly improves) the ratio for group *B* images while for all image groups its compression performance is better compared to RCT. Although applying color space transforms does not cause large improvements in terms of compression ratio, it should be noted that on the other hand these improvements are greater than differences in ratios between certain algorithms, e.g.

using JPEG-LS with RCT transform gives better ratios than CALIC for untransformed RGB images. We also note that the RCT transform is covered by the JPEG-LS extended standard (ISO/IEC and ITU-T 2002a).

The ratios for images after histogram packing or color space transforms and histogram packing are reported, for group *D* images only, in Table 5. Color channels of retinal images used for experiments in this paper are not very sparse, only in the case of group *D* images the sparseness is noticeable (averaged over 3 channels, below 95% levels are active). For other images the histogram packing did not change compression ratios noticeably. For all the algorithms the improvement in compression ratio due to histogram packing is about 3%. The improvement for images subjected to both color space conversion and histogram packing roughly equals the sum of individual improvements of histogram packing and color space conversion respectively.

Table 5: Compression ratios after histogram packing [bpp]

Color transform?	PNG	Lossless JPEG	JPEG-LS	JPEG2000*	JPEG2000	CALIC
no	7.136	8.248	7.205	7.474	7.474	7.046
yes	6.824	7.367	6.075	6.461	6.576	5.829

CONCLUSIONS

We analyzed the performance of several lossless image compression algorithms, namely Lossless JPEG, JPEG-LS, JPEG2000, PNG, and CALIC for a large set of medical retinal images. JPEG-LS was found to be the best performing algorithm when we considering both the compression speed and the achieved compression ratio. For nearly all the images it gives compression ratios close to the best performing algorithm (CALIC) while in terms of compression speed it slightly lacks behind Lossless JPEG, the fastest of the tested algorithms. Conveniently, apart being a standard itself, JPEG-LS is also incorporated in the DICOM standard and is hence readily available in PACS systems.

We also analyzed the effects of applying reversible transforms such as color space and histogram packing transforms, applied to the images prior to compression. We found that color space conversions are useful in the average case, although for certain images it might provide worse compression ratios. Since only some of images are of sparse histograms histogram packing has an influence on some images only.

REFERENCES

Adams, M.D. 2001. The JPEG-2000 Still Image Compression Standard. ISO JPEG Working Group Document ISO/IEC JTC 1/SC 29/WG 1 N 2412.

Christopoulos C.; A. Skodras and T. Ebrahimi. 2000. “The JPEG2000 Still Image Coding System an Overview”. *IEEE Transactions on Consumer Electronics* 46(4), 1103-27.

Consultative Committee for Space Data Systems. 1997. *Lossless Data Compression. CCSDS Recommendation for Space System Data Standards*. CCSDS 121.0-B-1, Blue Book.

Ferreira, P.J.S.G. and A.J. Pinho. 2002. “Why does histogram packing improve lossless compression rates?”. *IEEE Signal Processing Letters* 9(8), 259-61.

Golomb, S.W. 1966. “Run-Length Encodings”. *IEEE Trans. on Information Theory* IT-12, 399-401.

Huffman, D.A. 1952. “A method for the construction of minimum-redundancy codes”. *Proceedings of the Institute of Radio Engineers* 40(9), 1098-101.

ISO/IEC and ITU-T. 1999. *Information technology – Lossless and near-lossless compression of continuous-tone still images – Baseline.* ISO/IEC International Standard 14495-1 and ITU-T Recommendation T.87.

ISO/IEC and ITU-T. 2002a. *Information technology – Lossless and near-lossless compression of continuous-tone still images: Extensions.* ISO/IEC International Standard 14495-2 and ITU-T Recommendation T.870.

ISO/IEC and ITU-T. 2002b. *Information technology – JPEG 2000 image coding system: Core coding system.* ISO/IEC International Standard 15444-1 and ITU-T Recommendation T.800.

Langdon, G.; A. Gulati and E. Seiler. 1992. “On the JPEG model for lossless image compression”. In *Proceedings DCC '92, Data Compression Conference,* pp. 172-80.

Moffat, A.; R.M. Neal and I.H. Witten. 1998. “Arithmetic Coding Revisited”. *ACM Transactions on Information Systems* 16(3), 256-94.

National Electrical Manufacturers Association. 2001. *Digital Imaging and Communications in Medicine (DICOM) Part 5: Data Structures and Encoding.* NEMA PS 3.5-2001.

Pinho, A.J. 2001. “On the impact of histogram sparseness on some lossless image compression techniques”. In *Proc. of the IEEE Int. Conf. on Image Processing,* vol. III, pp. 442-5.

Rice, R.F. 1979. “Some practical universal noiseless coding techniques – part III”. Jet Propulsion Laboratory tech. report JPL-79-22.

Schaefer, G; R. Starosolski and S.Y. Zhu, S.Y. 2005. “An evaluation of lossless compression algorithms for medical infrared images”. In *Proc. of 27th IEEE Int. Conference Engineering in Medicine and Biology EMBC2005,* pp. 1125-28.

Starosolski, R. 2005. “Compressing images of sparse histograms”. In *Proc. SPIE, Medical Imaging,* SPIE Vol. 5959, pp. 209-17.

Starosolski, R. 2006. “Simple Fast and Adaptive Lossless Image Compression Algorithm”, to appear in *Software - Practice and Experience*.

Weinberger M.J.; G. Seroussi and G. Sapiro. 1996. “LOCO-I: A low complexity, context based, lossless image compression algorithm”. In *Proceedings DCC'96, Data Compression Conference,* pp. 140-9.

Weinberger M.J.; G. Seroussi and G. Sapiro. 2000. “The LOCO-I lossless image compression algorithm: Principles and standardization into JPEG-LS”. *IEEE Trans. Image Processing* 9(8), 1309-24.

Wu, X. 1997. “Efficient Lossless Compression of Continuous-tone Images via Context Selection and Quantization”. *IEEE Trans. on Image Processing* IP-6, 656-64.

Wu, X. and N. Memon. 1997. "Context-based, Adaptive, Lossless Image Codec". *IEEE Transactions on Communications* 45(4), 437-44.

WWW Consortium. 1996. *PNG (Portable Network Graphics) Specification, Version 1.0*. WWW Consortium Recommendation, http://www.w3.org/TR/REC-png.html.

Ziv J. and A. Lempel. 1977. "A universal algorithm for sequential data compression". *IEEE Transactions on Information Theory* 32(3), 337-43.

ACKNOWLEDGMENTS

This work was supported by the British Council and the Ministry of Education and Science of the Republic of Poland under grant number WAR/342/44.