Task-Oriented Lossy Compression of Magnetic Resonance Images

Mark Anderson

ISG Technologies, Inc.
6509 Airport Road, Mississauga, ON, Canada L4V 1S7

M. Stella Atkins

School of Computing Science, Simon Fraser University
Burnaby, BC, Canada V5A 1S6

Jacques Vaisey

School of Engineering Science, Simon Fraser University
Burnaby, BC, Canada V5A 1S6

ABSTRACT

A new task-oriented image quality metric is used to quantify the effects of distortion introduced into magnetic resonance images by lossy compression. This metric measures the similarity between a radiologist's manual segmentation of pathological features in the original images and the automated segmentations performed on the original and compressed images. The images are compressed using a general wavelet-based lossy image compression technique, embedded zerotree coding, and segmented using a three-dimensional stochastic model-based tissue segmentation algorithm. The performance of the compression system is then enhanced by compressing different regions of the image volume at different bit rates, guided by prior knowledge about the location of important anatomical regions in the image. Application of the new system to magnetic resonance images is shown to produce compression results superior to the conventional methods, both subjectively and with respect to the segmentation similarity metric.

Keywords: magnetic resonance image, lossy compression, tissue segmentation, zerotree, wavelet, similarity index.

1 Introduction

Image compression can help to reduce the storage space and transmission time requirements of sizable medical images. The amount of compression that can be achieved using lossless methods may be inadequate for the large quantities of data generated in a medical imaging environment.[1] Lossy

compression techniques, which can achieve higher compression ratios at the expense of reconstructed image quality, have not been adopted in the medical imaging community due to the perceived or actual distortion of clinically significant image detail. However, it may be beneficial to employ lossy compression techniques for purposes other than diagnosis, such as archival storage. Similarly, when a large image is being retrieved from storage or transmitted to a remote location, it may be beneficial to first transmit a lower-quality approximation for initial examination, followed by the slower transmission of further detail, if necessary.

When lossy compression is used, it is desirable to preserve as much clinically useful information as possible. The nature of the information that must be retained and that may be discarded depends on both the type of image and the diagnostic task to be performed. There exist a number of objective measures of reconstructed image quality and distortion, but the relationship of these measures to the actual performance of applications on the reconstructed images has not been clearly identified. In this paper, the effects of certain types of lossy compression on a particular medical imaging application are investigated; a new method for the objective quantification of loss in these medical images with respect to the application is proposed; and a new system is presented for the lossy compression of medical images that takes advantage of certain known characteristics of the data and of the applications that will be applied to the data, with the goal of improving compression performance while maintaining the effectiveness of the applied medical imaging task.

In Section 2, wavelet-based image compression and a state-of-the-art wavelet-based still-image compression technique, known as embedded zerotrees of wavelet coefficients (or EZW) coding is described in general terms. In Section 3, a particular medical imaging application, namely the segmentation of white-matter brain lesions in magnetic resonance volume data, is described. A quality measure based on the similarity between such image segmentations is defined. In Section 4, two new region-based variants of EZW are presented, and in Section 5, results of conventional and region-based EZW coding of test images are presented, using the segmentation similarity as a quality metric. The conclusions are summarized in Section 6.

2 Zerotree Coding

The wavelet transform is a hierarchical subband decomposition particularly suited to image compression. Image compression using wavelet transform techniques is concerned with efficiently quantizing and coding the coefficients generated by the wavelet transform and which represent the portions of the image energy at various scales and orientations. Shapiro[2–4] introduced an image coder based on the *zerotree*, a data structure which efficiently represents the low-magnitude wavelet coefficients. Shapiro's technique is known as embedded zerotrees of wavelet coefficients, or EZW, coding.

In EZW coding, the significant wavelet coefficients are gradually identified and their precisions refined, roughly in order of their magnitudes. The encoding or decoding may be terminated at any point, and the resulting bit stream is a prefix of all lower-rate encodings; this is referred to as *embedded coding*. Thus, compression is achieved by terminating the transmission or storage of the embedded code at some point in the bit stream, and the exact bit rate is controlled by choosing the point at which this termination takes place. Embedded coding schemes naturally suit a progressive mode of transmission of the image: at any time, images reconstructed from the decoded bit stream can be displayed as

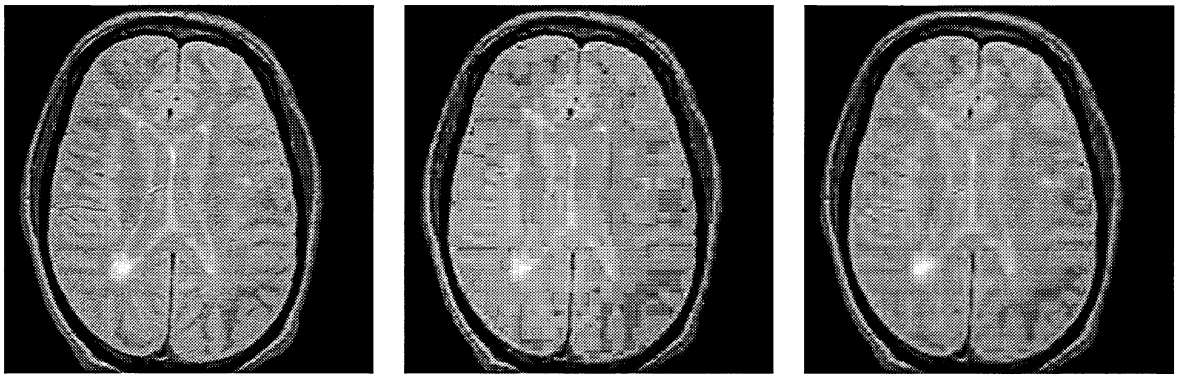

(a) (b) (c)

Figure 1: (a) The original MR image slice. (b) The JPEG-coded image reconstructed at 0.25 bpp. (c) The EZW′-coded image reconstructed at 0.25 bpp.

increasingly good approximations of the source image.

We have implemented an image coder and decoder based on EZW, but with a few simplifications that result in a slight decrease in coding performance compared to Shapiro's EZW implementation. In this paper, our codec will be referred to as EZW′ in order to distinguish it from Shapiro's EZW. Refer to Anderson[5] for a complete description of EZW′.

Our codec was implemented in the WiT ™ visual data-flow image processing environment.[6,7] The two-dimensional wavelet transform was implemented using a modified version of the public-domain `wvlt` library available from the Imager group at the University of British Columbia.[8] As in Shapiro's work, the entropy coder was an arithmetic coder based on the one described by Witten et al.[9] Shapiro used a set of quadrature mirror filters described by Adelson and Simoncelli.[10] The so-called pseudo-Coiflet filters[11] offer similar performance, and the 4-tap version was used for this work.

Figure 1 shows an MR image slice and its reconstructions using EZW′and the widely used JPEG lossy image compression algorithm,[12] both at about the same bit rate of 0.25 bits per pixel (bpp). In the JPEG image, the blocky reconstruction artifacts are conspicuous. In the EZW′-compressed image, loss of image detail is visible in the smooth central parts of the image; this blurring increases at lower bit rates. However, the quality of the EZW′-coded images is visually superior to the JPEG reconstructions at similar bit rates. The JPEG and EZW′ reconstructions are similar at rates down to about 0.5 bpp; however, at lower bit rates the EZW′ reconstruction quality degrades much more gracefully than JPEG.

3 Image Segmentation

We are interested in how compression results may be quantified and improved with respect to a particular medical image processing task: namely, segmentation of white-matter lesions present in brain MR images of multiple sclerosis patients. The segmentation technique, developed by Johnston et al.,[13–15] is used to semi-automatically segment brain tissues, and in particular MS lesions, in MR

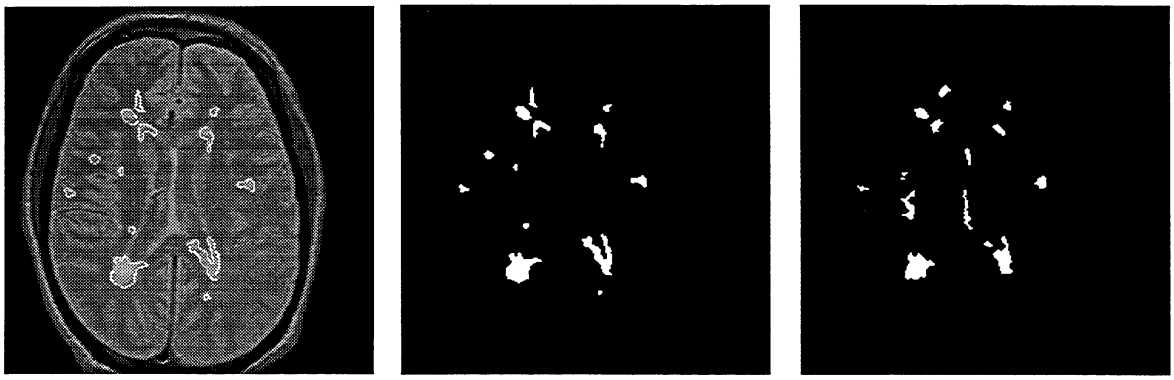

(a) (b) (c)

Figure 2: Lesion segmentations of an MR image slice. (a) The PD-weighted image overlaid with manually drawn lesion outlines. (b) The manual lesion segmentation as a binary mask. (c) The semi-automatic lesion segmentation. The similarity index calculated between (b) and (c) is 0.63.

images. This segmentation results in a classification of each voxel in the volume as "lesion" or "not lesion." Such a classification can be represented by a binary image, as depicted in Figure 2 for a single slice in our sample volume.

Xuan et al.[16] used the segmentation of compressed medical images as a method of quantifying the effects of compression loss on a typical medical image processing task. They compared the binary segmentations of original and compressed images, and used a percentage of misclassified pixels as the metric. A more sophisticated measure is the *similarity index* defined by Zijdenbos et al.[17,18] Consider a binary segmentation as a set A containing the pixels considered to belong to the classification. Then the similarity of two segmentations A_1 and A_2 is given by a real number $S \in \{0..1\}$ defined by

$$S = 2\frac{|A_1 \cap A_2|}{|A_1| + |A_2|}$$

This provides a simple numerical measure by which the quality of segmentations can be compared; for instance, the similarity index of the two binary images in Figure 2 (b) and (c) is 0.63. This definition can be applied equally to three-dimensional segmentations. For an objective image quality metric, we will calculate the similarity of the automated segmentation of the volume data with the gold standard provided by the radiologist. In this paper, all similarity index measurements are made with respect to the radiologist's manually outlined lesion segmentation.

4 Region-Based Compression

Various spatial regions of our image volumes are more "important" than others, for medical imaging applications such as brain tissue segmentation. It is useful, both for image intensity correction and for segmentation, to isolate the brain in each slice from the surrounding tissues such as skull, muscle and fat. Conventionally, this is performed manually,[19] but recently Mackiewich[20] has obtained very good results using automatic techniques. Thus, we have available a simple partition of the 3-dimensional

volume image into two disjoint images, namely the regions inside and outside the brain contour.

To improve compression, our task-oriented compression technique takes advantage of the brain contour information by compressing the two areas of the image (inside and outside the contour) at different bit rates, and by devoting more of the compressed data stream to those slices with larger image areas within the contour. To compress the image region inside the brain contour (the *interior* region) at a different rate from the *exterior* region, we partition the image and produce two compressed bit streams, which may be stored or transmitted together. For transmission applications, the bit streams could be multiplexed, while for archival applications, the streams could be stored in separate files.

The straightforward way to split and code the separate regions is to partition pixels of the raw image in the spatial domain, producing two images: one containing all the pixels inside the brain contour (and all zero outside the contour) and the second containing all the pixels outside the contour (and all zero inside). These images are then compressed separately using EZW′, producing the interior and exterior bit streams. We will call this technique spatial region-based (SRB) EZW′.

Our second partitioning technique takes advantage of the fact that each wavelet coefficient represents both spatial and frequency information. This spatial localization allows us to identify coefficients in each subband corresponding to the spatial regions we have identified in the original volume, thereby permitting application of the partitioning mask in the wavelet domain. First, the wavelet transform is applied to the raw image. The brain mask is modified to cover the spatially corresponding coefficients in each subband, and the coefficient image is partitioned using the new mask. Finally, each coefficient partition is coded individually with EZW′, again producing the two bit streams. This technique will be called transform-domain region-based (TRB) EZW′.

In both SRB-EZW′and TRB-EZW′, the interior and exterior images are coded independently, using parameters and initial thresholds calculated from the two individual partitioned images. The two decoded images are recombined by summing them, which can be seen as special case of the wavelet-based method for image fusion described by Li et al..[21]

The EZW′ codec provides for fine control of the bit rate achieved by specifying the number of bits emitted or consumed. Therefore, to code the image at a given overall bit rate using the region-based EZW′ coders, it is necessary first to determine the number of bits with which to code the interior and exterior regions. We can define a real-valued parameter $p \in [0, 1]$, which we will call the *importance ratio* of the interior region with respect to the whole image, as the ratio of the number of bits used to code the region to the number of bits used to code the whole image.

5 Results

5.1 Conventional EZW′

The MR data set used for this work was acquired for the MS/MRI Study Group at the University of British Columbia.[22] It is a dual echo MRI sequence of a brain consisting of 27 slices of 256×256 voxels per slice, with a slice thickness of 5 mm and no inter-slice gap. The raw data was stored at 16 bpp, but information was contained in only 12 of the 16 bits; furthermore, the data was scaled linearly

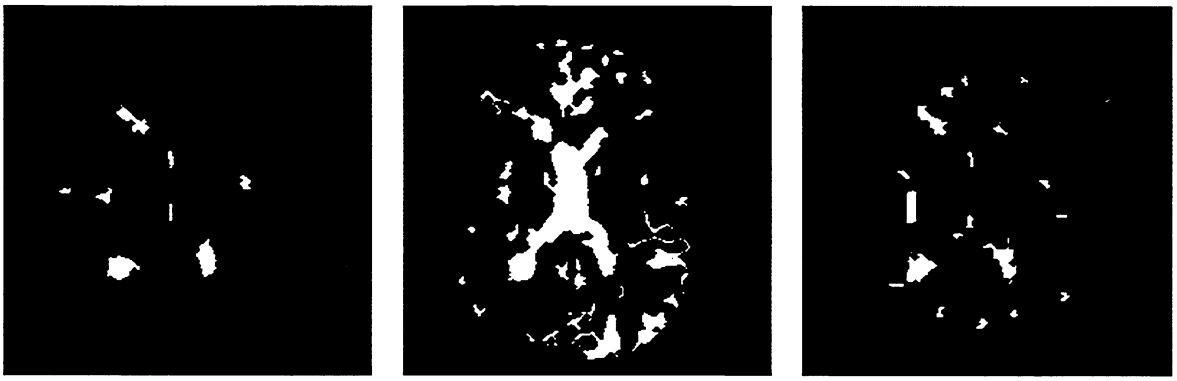

(a) (b) (c)

Figure 3: Segmentations of the EZW′-compressed data set reconstructed (a) at 0.25 bpp, and (b) at 0.125 bpp. Slice similarity indices are (a) 0.60 and (b) 0.20. For comparison, segmentation of a slice compressed at 0.25 bpp using JPEG is shown in (c); this segmentation has similarity index 0.52.

to 8 bpp and supplied in that form. Thus the original data set occupied 6.75 MB of storage space; the scaled data occupied 3.375 MB.

The scaled data was segmented using the semi-automatic method discussed in Section 3. The final lesion segmentation in each case was performed only for slices 11 through 20, which contain the majority of the lesion tissue. The aggregate similarity index for this ten-slice volume is 0.51.

The image volume was compressed slice-by-slice, using the EZW′ implementation described in Section 4. Using the compression ratios and image qualities described in Shapiro's paper[3] as a guide, we chose to reconstruct the slices at 0.25 bpp and at 0.125 bpp. The reconstructed image volumes were preprocessed and segmented as described above. The segmentation of one slice is shown in Figure 3. Compare these segmentations to those given in Figure 2 for the same slice on the raw (not compressed) data. The lesion segmentation obtained using the 0.25 bpp EZW′-compressed data visually resembles both the manual and the automatically generated raw-data segmentations, though some detected lesions have changed in shape, and others have disappeared altogether. The segmentation of the 0.125 bpp data is very distorted, with a large amount of falsely identified lesion. The degree of distortion seen in the images is reflected in their similarity indices of 0.60 and 0.20 respectively. For comparison, the segmentation of the same slice compressed at about 0.27 bpp using JPEG is shown in Figure 3 (c). This segmentation, with a similarity index of 0.52, is slightly worse than the EZW′ segmentation similarity of 0.60 at the same bit rate.

Similar results are found using 3D volume images instead of single slices. The similarity indices for the ten-slice volume segmentations are given in Table 1. The segmentation similarity does not suffer greatly as a result of the EZW′ compression at 0.25 bpp; here the volume similarity index decreases by about 14 per cent. The JPEG compression at a slightly higher bit rate results in a reduction of the volume similarity index by about 20 per cent. EZW′ compression at 0.125 bpp, however, results in a similarity reduction of about 74 per cent. Note that the compressed bit rate could be further reduced by implementing EZW in its entirety.

Table 1: Aggregate similarity indices for slices 11 through 20 of the automatic lesion segmentations of raw and compressed data.

Data Set	Similarity
Raw 8 bpp	0.51
EZW′ 0.25 bpp	0.44
JPEG 0.27 bpp	0.40
EZW′ 0.125 bpp	0.13

5.2 Region-Based EZW′

Each slice of our test MRI volume (both PD and T_2 weighted) was compressed using region-based EZW′(both the SRB and TRB variants), using a manually generated brain mask, though equally good or better masks recently have been produced automatically.[20] To simulate results at various bit rates, we first compressed the interior and exterior data of each slice separately at 1 bpp using the region-based coders. Due to the embedded nature of the EZW encoding, we were then able to reconstruct images from these individual bit-stream files at any rate less than or equal to 1 bpp, in order to simulate coding and decoding or transmission at various rates.

We have seen that a range of bit rates are possible for the interior and exterior subimages while maintaining a fixed overall bit rate. The importance ratio parameter p can be varied between 0 and 1 to jointly control the quality of the two subimages; larger values of p result in an improvement in the quality of the interior subimage at the expense of the exterior subimage. The interior and exterior bit rates resulting from the use of various p values for the test volume are given in Table 2.

Table 2: Interior and exterior bit rates used to reconstruct the sample slice with region-based EZW′ at 0.25 bpp overall.

	Bit Rate (bpp)	
p	Interior	Exterior
1.0	0.53	0
0.9	0.48	0.047
0.7	0.37	0.14

The effects of region-based EZW′ reconstruction on image quality at an overall bit rate of 0.25 bpp using various values of p are depicted in Figure 4 (SRB-EZW′) and Figure 5 (TRB-EZW′). Visually, no improvement in image quality of the brain region over conventional EZW′ is apparent in the SRB-EZW′ reconstructions, for any of the values of p chosen; when using TRB-EZW′, however, increased detail in the central brain region is visible in the images reconstructed using $p = 0.9$ and $p = 0.7$. Restricting our attention to the quality of fine detail in the interior region, the image quality of the TRB-EZW′ reconstructions is superior to SRB-EZW′ in each case. When $p = 0.5$, the image quality of TRB-EZW′ is similar to that obtained with conventional EZW′.

Both SRB-EZW′ and TRB-EZW′ introduce artifacts in the reconstructed image near the location of the brain mask edge. In SRB-EZW′, the artifacts appear as irregular bright ridges around the brain

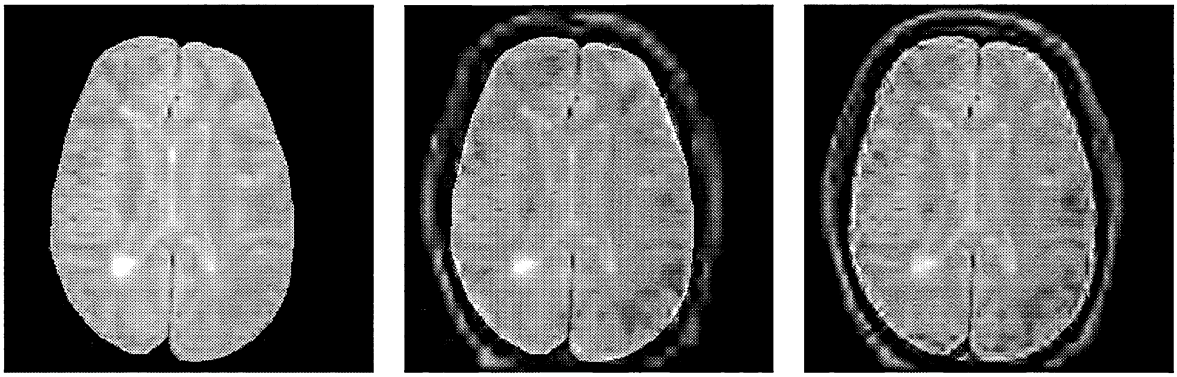

(a) (b) (c)

Figure 4: Results of SRB-EZW′ coding of the test slice, with (a) $p = 1.0$, (b) $p = 0.9$, and (c) $p = 0.7$.

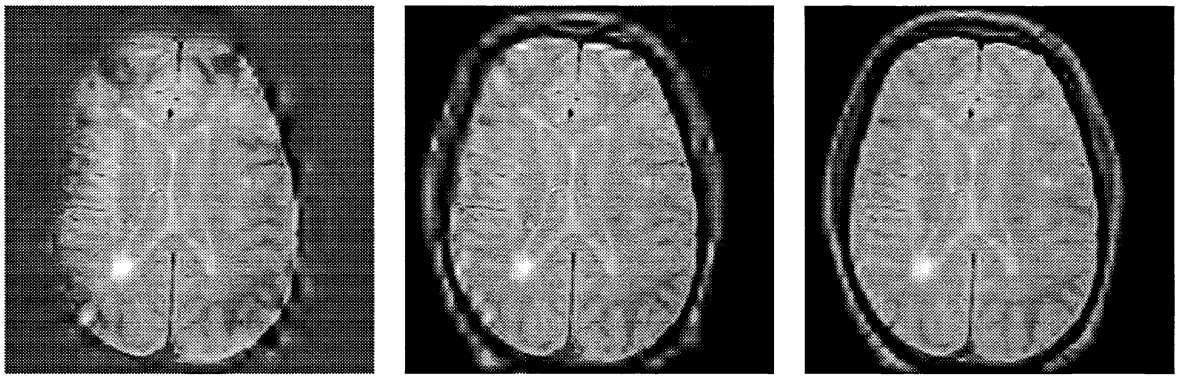

(a) (b) (c)

Figure 5: Results of TRB-EZW′ coding of the test slice, with (a) $p = 1.0$, (b) $p = 0.9$, and (c) $p = 0.7$.

where the two subimages meet. These artifacts are caused by the imperfect edge reconstruction inherent in lossy wavelet compression. By masking the brain in the spatial domain, sharp artificial edges are introduced in both the interior and exterior subimages, resulting in artifacts in both reconstructions; when these subimages are summed to obtain the final image, the visible artifacts are emphasized.

In SRB-EZW′, the sharp artificial edges in the partitioned images can result in corresponding large-valued coefficients in the wavelet transform domain. Since EZW′ codes the largest coefficients first, the algorithm expends more bits coding these edges at the expense of smoother regions. TRB-EZW′, by partitioning the image in the transform domain, avoids the introduction of artificial edges in the data, and is able to reconstruct the interior subimage with more detail.

For all values of p, the SRB-EZW′ artifacts have a similar appearance and are spatially quite localized. In TRB-EZW′, however, the artifacts are very pronounced for large values of p, but they become less prominent as the value of p decreases. The TRB-EZW′ artifacts are due to the the distortion present in the exterior subimage when it is reconstructed at a very low bit rate.

The entire manually-generated 3D binary brain mask contains 550,315 voxels, representing 31.1 per

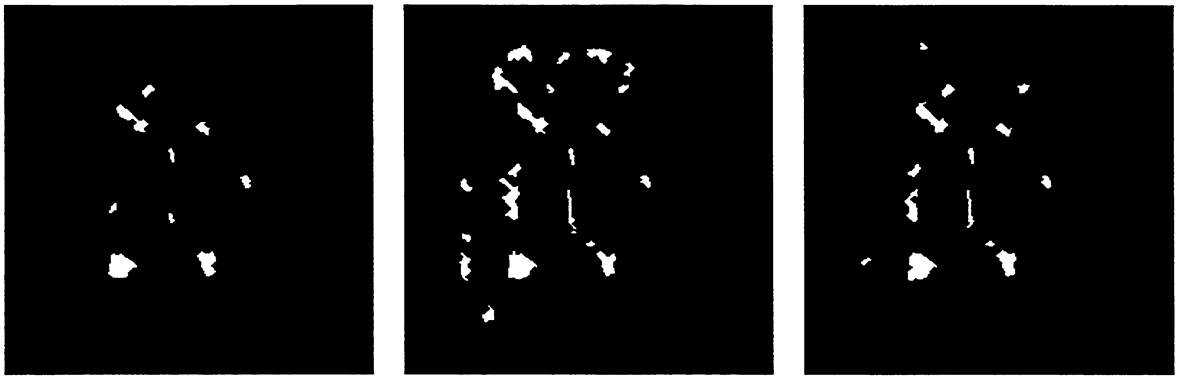

(a) (b) (c)

Figure 6: Semi-automatic segmentations of data compressed to 0.25 bpp using (a) SRB-EZW′ with $p = 1.0$, and TRB-EZW′ with (b) $p = 0.9$ and (c) $p = 0.7$.

cent of the voxels in the whole volume. For comparison with our standard EZW′ compressed data set, we reconstructed the volume at 0.25 bpp in total using TRB-EZW′ with 90 and 70 per cent of the total bit budget allocated to the interior voxels (i.e. with importance ratios of $p = 0.9$ and $p = 0.7$). The volume was also reconstructed at 0.25 bpp using SRB-EZW′; in this case, an importance ratio of $p = 1.0$ was used, since lower values of p resulted in worse image quality than TRB-EZW′. The corresponding interior and exterior bit rates are given in Table 3.

Table 3: Interior and exterior bit rates used to reconstruct the ten-slice volume with region-based EZW′ at 0.25 bpp overall.

	Bit Rate (bpp)	
p	Interior	Exterior
1.0	0.80	0
0.9	0.72	0.036
0.7	0.56	0.11

The automatic segmentation procedure was applied to the volume data sets reconstructed using SRB-EZW′ and TRB-EZW′. Figure 6 (a) depicts the segmentation of data compressed to 0.25 bpp using SRB-EZW′ with $p = 1.0$. This segmentation classifies most of the large lesions well; many of the smaller lesion areas detected in the segmentation of the raw data are missed here, but these smaller lesions do not greatly affect the similarity measure. Figure 6 (b) and (c) show the segmentations of the same slice reconstructed using TRB-EZW′ with $p = 0.9$ and $p = 0.7$, respectively. Although the lesion areas inside the brain have again been classified quite well, large areas of bright edge artifacts have been misclassified as lesion, especially in the volume reconstructed using $p = 0.9$. When $p = 0.7$, these false positives are considerably reduced.

The results of the similarity comparisons with the radiologist's segmentation are shown in Table 4. Each form of compression at 0.25 bpp resulted in a decrease in the segmentation similarity index. The largest drop, about 30 per cent, was associated with compression using TRB-EZW′ with $p = 0.9$,

which was shown to produce significant edge artifacts in the reconstructed images. Both TRB-EZW′ compression with $p = 0.9$ and SRB-EZW′ compression with $p = 1$ resulted in the smallest similarity index decrease of only about 8 per cent; however, the former method retained lower-quality exterior image data while the latter did not. The conventional EZW′ compression gave similar but slightly worse results.

Further experiments with additional data sets are required to determine the significance of the small differences in similarity index we have reported. However, our results indicate that the region-based methods are able to maintain or improve the suitability of reconstructed images for our task while simultaneously allowing for variable-quality reconstruction of different image regions.

Table 4: Aggregate similarity indices for slices 11 through 20 of the automatic lesion segmentations of raw and compressed data.

Data Set	Similarity
Raw 8 bpp	0.51
EZW′ 0.25 bpp	0.44
SRB-EZW′ 0.25 bpp, $p = 1$	0.47
TRB-EZW′ 0.25 bpp, $p = 0.9$	0.35
TRB-EZW′ 0.25 bpp, $p = 0.7$	0.47

6 Conclusions

In this paper, we have introduced a new objective task-oriented image-quality metric which can be used to judge the degree of degradation of images which have been reconstructed using lossy compression techniques. The metric, which uses segmentation as its medical image processing application, calculates the similarity between the results of an automatic segmentation and a gold standard segmentation such as that generated manually by a radiologist.

We have adapted a general, state of the art still-image compression system, EZW′, for use with three-dimensional MR image volumes, exploiting characteristics of both the data being compressed and the task that is applied to the data by incorporating information about the spatial location of important regions in the three-dimensional volume to to improve the reconstruction quality of these regions at the expense of the less important regions. We have implemented two techniques for applying this information to the image compression process, one in the spatial domain and one in the transform domain.

The results of using the region-based EZW′ compression techniques on a sample MRI volume were compared with the conventional EZW′ and with standard JPEG image compression results. We have found that our TRB-EZW′ method, which performs the image partition in the wavelet-transform domain, provides the best image quality measured both subjectively and using the segmentation similarity metric, while still providing variable-quality reconstruction of the image regions.

7 Acknowledgments

The authors wish to thank Don Paty, Andrew Riddehough, Keith Cover and Brenda Rhodes of the University of British Columbia MS/MRI Study Group for supplying the MRI data and the manually drawn lesion outlines on which the algorithm was tested. Thanks also to Bob Lewis of the UBC Imager Lab for his help with wavelet filters. This research has been partially funded by the Natural Sciences and Engineering Research Council of Canada.

8 REFERENCES

[1] Pamela A. Cosman, Robert M. Gray, and Richard A. Olshen. Evaluating quality of compressed medical images: SNR, subjective rating, and diagnostic accuracy. *Proceedings of the IEEE*, 82(6):919–932, June 1994.

[2] Jerome M. Shapiro. An embedded wavelet hierarchical image coder. In *Proceedings of ICASSP'92*, volume IV, pages 657–660. IEEE, 1992.

[3] Jerome M. Shapiro. Embedded image coding using zerotrees of wavelet coefficients. *IEEE Transactions on Signal Processing*, 41(12):3445–3462, December 1993.

[4] Jerome M. Shapiro. An embedded hierarchical image coder using zerotrees of wavelet coefficients. In James A. Storer and Martin Cohn, editors, *DCC '93 : Data Compression Conference*, pages 214–223. IEEE, April 1993.

[5] Mark C. Anderson. Task-oriented lossy compression of magnetic resonance images. M.Sc. thesis, Simon Fraser University, August 1995.

[6] T. Arden and J. Poon. *WIT User's Guide.* Logical Vision Ltd., Burnaby, BC, October 1993. Version 4.1.

[7] Chitra Balasubramaniam. Dataflow image processing. *Computer*, 27(11):81–84, November 1994.

[8] Bob Lewis. UBC Imager wavelet library (`wvlt`). Department of Computer Science, University of British Columbia. Release 1.3, September 1994.

[9] Ian H. Witten, Radford M. Neal, and John G. Cleary. Arithmetic coding for data compression. *Communications of the ACM*, 30(6):520–540, June 1987.

[10] Edward H. Adelson and Eero Simoncelli. Orthogonal pyramid transforms for image coding. In T. Russell Hsing, editor, *Visual Communications and Image Processing II*, Proceedings of SPIE 845, pages 50–58. SPIE, 1987.

[11] Leena-Maija Reissell. Multiresolution geometric algorithms using wavelets I: representation for parametric curves and surfaces. Technical Report 93-17, University of British Columbia, Dept. of Computer Science, May 1993.

[12] William B. Pennebaker and Joan L. Mitchell. *JPEG Still Image Data Compression Standard.* Van Nostrand Reinhold, New York, 1993.

[13] Brian G. Johnston. Three-dimensional multispectral stochastic image segmentation. M.Sc. thesis, University of British Columbia, January 1994.

[14] Brian G. Johnston, M. Stella Atkins, and Kellogg S. Booth. Partial volume segmentation in 3D of lesions and tissues in magnetic resonance images. In Murray H. Loew, editor, *Medical Imaging 1994: Image Processing*, Proceedings of SPIE 2167, pages 28–39. SPIE, 1994.

[15] Brian G. Johnston, M. Stella Atkins, Blair T. Mackiewich, and Mark C. Anderson. Segmentation of multiple sclerosis lesions in intensity corrected multispectral MRI. *IEEE Transactions on Medical Imaging*, April 1996. To appear.

[16] Jianhua Xuan, Tülay Adah, Yue Wang, and Richard Steinman. Predictive tree-structured vector quantization for medical image compression and its evaluation with computerized image analysis. In Yongmin Kim, editor, *Medical Imaging 1995: Image Display*, Proceedings of SPIE 2431. SPIE, 1995.

[17] Alex P. Zijdenbos, Benoit M. Dawant, Richard A. Margolin, and Andrew C. Palmer. Morphometric analysis of white matter lesions in MR images: method and validation. *IEEE Transactions on Medical Imaging*, 13(4):716–724, December 1994.

[18] Alex P. Zijdenbos. *MRI Segmentation and the Quantification of White Matter Lesions.* Ph.D. thesis, Vanderbilt University, December 1994.

[19] Thad Q. Bartlett, Michael W. Vannier, Daniel W. McKeel, Jr., Mokhtar Gado, Charles F. Hildebolt, and Ronald Walkup. Interactive segmentation of cerebral gray matter, white matter, and CSF: photographic and MR images. *Computerized Medical Imaging and Graphics*, 18(6):449–460, 1994.

[20] Blair T. Mackiewich. Intracranial boundary detection and radio frequency correction in magnetic resonance images. M.Sc. thesis, Simon Fraser University, August 1995.

[21] H. Li, B. S. Manjunath, and S. K. Mitra. Multisensor image fusion using the wavelet transform. *Graphical Models and Image Processing*, 57(3):235–245, May 1995.

[22] D. W. Paty, D. K. B. Li, the UBC MS/MRI Study Group, and the IFNB Multiple Sclerosis Study Group. Interferon beta-1b is effective in relapsing-remitting multiple sclerosis. II. MRI analysis results of a multicenter, randomized, double-blind placebo-controlled trial. *Neurology*, 43(4):662–667, April 1993.